Analysing Observables in Structure Formation Theories

Rita Tojeiro

Institute for Astronomy
School of Physics

University of Edinburgh
Doctor of Philosophy

September 2008

Abstract

This thesis presents an exploration of various aspects relating to the formation and evolution of structure in the Universe. It focuses on two main observables which provide information on two distinct epochs of the Universe: Part I analyses the Cosmic Microwave Background (CMB) which is used to test early Universe theories and validate current methods for cosmological parameters estimation; Part II analyses the distribution, history and content of local galaxies with a view to learn about type Ia supernovae progenitors, assembly of stellar mass in galaxies and galaxy evolution.

In Part I, a search for signs of non-Gaussianity in the Wilkinson Microwave Anisotropy Probe is conducted, using the two-point correlation function of peaks (hot and cold spots) in the temperature field. A clear deviation from Gaussianity is found in both data releases, which is associated with cold spots, the southern hemisphere, large-scales and the galactic plane. The results indicate that the presence of un-subtracted foregrounds in the data are a more likely explanation for this signal than a cosmological origin, but the latter cannot be excluded. Part I further explores the two-point correlation function of temperature peaks as an estimator to constrain f_{NL}, a specific type of non-Gaussianity. Using sets of non-Gaussian simulated maps with the correct cosmology and resolution, this thesis explores how accurately one can hope to constrain f_{NL} when data from the upcoming CMB experiment Planck is available.

Part II presents a novel method developed to extract the star formation history of a galaxy from its spectrum: VErsatile SPectral Analysis (VESPA). VESPA dynamically adapts the number of parameters it recovers from each spectrum to each galaxy, only recovering as much information as the data warrant. This insures the recovered solutions are dominated by the signal, not the noise, and allows robust recovery of star formation and metallicity histories and up to two dust extinction values per galaxy. VESPA was applied to the fifth data release of the Sloan Digital Sky Survey (SDSS)

to construct a catalogue of histories for nearly half a million galaxies. Part II also explores how this catalogue can provide important information about the formation of structure in the local Universe. Specifically, this thesis: shows evidence for the presence of a short-lived progenitor to SNIa and explores VESPA's potential role in future Dark Energy SNIa experiments; explores how changes in the Initial Mass Function affect stellar mass estimates and its consequences; and by using estimates of the galaxy mass function as a function of redshift (calculated using both the spectral fossil record and instantaneous star formation rate methods) explores the possibility of putting constraints on the merger history of present-day galaxies.

Declaration

I declare that this thesis is not substantially the same as any that I have submitted for a degree or diploma or other qualification at any other University. I further state that no part of my thesis has already been or is being concurrently submitted for any such degree, diploma or other qualification.

This thesis is the outcome of my own work except where specifically indicated in the text. Some of the results in this thesis were done in collaborative work, and my contribution to each project is as follows.

- The research in Chapters 3 and 4 was done in collaboration with the authors in Tojeiro et al 2006 and 2007, respectively; I made significant contributions to the development of the methods and was responsible entirely for the implementation.
- Section 5.3.4 was done in collaboration with Stephen Wilkins, Andrew Hopkins and Neil Trentham. My main contribution consisted of adapting the VESPA algorithm in order to produce a dataset which was necessary for the project.
- Chapter 6 was done in collaboration with Eric Aubourg, Raul Jimenez, Alan Heavens, Michael Strauss and David Spergel. In addition to significant contributions to the development of the project and general discussion, I contributed by supplying the star formation histories data.

Part of this thesis is based on the following published or submitted scientific articles:

- R. Tojeiro, P.G. Castro, A. F. Heavens, S. Gupta, 2006,MNRAS, 365, 265.
- R. Tojeiro, A.F. Heavens, R. Jimenez, B. Panter, 2007, MNRAS, 381, 1252.
- E. Aubourg, R. Tojeiro, R. Jimenez, A. F. Heavens, M. Strauss, D. Spergel, 2008, submitted.

- S. M. Wilkins, A. M. Hopkins, N. Trentham, R. Tojeiro, 2008, MNRAS, in print.

Rita Tojeiro
Edinburgh,
September 2008.

Acknowledgements

For a long time it really didn't feel like I'd get to write this part, but now here we are. So for those students who, like me, like to look through strangers' theses and browse the acknowledgements section here's a message of hope: don't panic.

On a more serious note, this thesis would have never been possible without the help and support of a number of nice people to whom I'd like to send my thanks.

My first thank you goes to the two people who have always supported my crazy wish to pursue a career in Astronomy, even when that meant leaving home and moving to a cold and slightly strange country up in north of the world - Pai, Mãe, um muito sincero obrigada.

Then to Alan, for his guidance, insight, patience and encouragement and for being all the things a good supervisor should be. Thank you for seeing through my insecurities and gently pushing me into the postdoc world, well beyond my comfort zone. These are early days, but you know, so far so good..!

To Ben and Patricia, for answering my questions with a straight face. Thank you for always being there and for constantly taking the time to give constructive and relevant criticism. I'm not any smarter for being a doctor, so you should expect a lot more coming your way - let me also thank you in advance for that.

To all my science collaborators, for their ideas and encouragement. Particularly I want to thank Raul for pushing this PhD forward at times when I really didn't have the energy to, and for being relentless in his task of finding me more and more things to do. Also to Andy and Paul, for their time in reading this thesis and for what was a nearly-enjoyable viva experience, and a much more insightful and much less frightening event than I ever thought possible.

To all the people who shared this PhD with me in the pub, at home, and in the office - above all I wouldn't have made it without you (or the beer). To Liv for making a flat a home; to Stuart for always making me see the beautiful side of the Universe; to Caz for turning the office into somewhere warm and friendly; to Michael for being ever-present; to Emma, Niall, Miller, Matt, Tom(s), Anita, Barney and Rachel for the golden afternoon-coffee times which I'll always warmly remember. To Berian for being curious. To Sam for making me laugh. To everyone who has ever bought me a pint (and Eric comes high on this list).

To Joana, Miana and Sónia for what is truly an eternal friendship, and to Vanessa for endless support, friendship, and the comfort of knowing there is someone who will always go out for a Francesinha with me. To the rest of my family who, in spite of the distance, managed to stay close and supportive, and are the reason why every trip back to Portugal is always a genuine pleasure.

To Dan, Dave, Jenny and Tania for giving me the most wonderful opportunity of sharing my science with the public - I couldn't think of a better team to do this with. I often think I could have easily packed up and left if it hadn't been for the constant kick of enthusiasm and focus that doing outreach allowed. To everyone who ever went on a Dark Sky Scotland weekend with me, and to everyone out there who has ever asked me a question I didn't know the answer to.

To those who helped without knowing! These are too many to list, but special mentions go to the open source community for the endless applications and snippets of code which made this PhD a far easier journey; to Beardo and Spinner for a podcast which turned a lot of unbearable early-morning and late-night moments into something enjoyable; and to the Bruichladdich distillery for making Bruichladdich.

Lastly but clearly not least, to Simon. For the love, the friendship and the huggin' - you are my world.

Contents

1 Introduction **1**

1.1 The smooth Universe 1

1.1.1 The geometry of the Universe 2

1.1.2 The dynamic Universe 2

1.1.3 The equation of state 4

1.2 The components of the Universe 5

1.2.1 Radiation 5

1.2.2 Baryonic matter 6

1.2.3 Dark Matter 7

1.2.4 Dark Energy 7

1.3 Observational cosmology tools 8

1.4 Structure formation 11

1.4.1 Inflation 11

1.4.2 The linear regime 14

1.4.3 The matter power-spectrum 17

1.4.4 The hierarchical model 18

1.4.5 Beyond the linear regime 20

1.4.6 Galaxy formation models 23

1.5 The Cosmic Microwave Background as an observable 24

1.5.1 The CMB observables 25

1.5.2 Relating angular sizes with linear scales 28

1.5.3 Physical mechanisms: the origin of the anisotropies 29

1.6 The integrated spectrum of a galaxy as an observable 35

1.6.1 Stellar population models 36

1.6.2 Dust models 37

1.6.3 Extracting the information 38

1.7 Summary 39

I Non-Gaussianity in the Cosmic Microwave Background 41

2 Background and methodology 43

2.1 Background 43
2.1.1 Searches for non-Gaussianity 44
2.1.2 f_{NL} models 45
2.1.3 Interpretation 46
2.2 Data 47
2.2.1 Foreground emission 48
2.2.2 Instrumental systematic effects 50
2.2.3 Estimating the temperature power-spectrum 51
2.2.4 CMB maps 53
2.3 Methodology 53
2.3.1 HEALPix 54
2.3.2 The Gaussian maps 54
2.3.3 High-pass filtered maps 56
2.3.4 Non-Gaussian maps 56
2.3.5 Estimating $\xi(\theta)$ 57
2.3.6 Statistics 58

3 Results 61

3.1 Year one 61
3.1.1 All-sky analysis 62
3.1.2 North-South analysis 63
3.1.3 Constraining in harmonic space 65
3.1.4 Constraining in real space 67
3.1.5 A single-frequency analysis 69
3.1.6 Removing the cosmological signal 70
3.1.7 Summary 73
3.1.8 Conclusions on the first year analysis 73
3.2 Year five 76
3.2.1 Full-sky analysis 77
3.2.2 North-South analysis 80

3.2.3 Constraining in real space . . . 83
3.2.4 The integrated Sachs-Wolfe effect . . . 87
3.2.5 Summary . . . 89
3.2.6 Conclusions on the fifth-year analysis . . . 89
3.3 f_{NL} constraints . . . 92
3.3.1 Summary . . . 94

II Exploring the local Universe with VESPA 97

4 VESPA 99
4.1 Background . . . 99
4.1.1 The cosmic star formation history . . . 100
4.1.2 Downsizing . . . 100
4.1.3 Stellar mass assembly . . . 103
4.1.4 The fossil record . . . 104
4.1.5 VESPA . . . 106
4.2 Method . . . 107
4.2.1 The problem . . . 107
4.2.2 The solution . . . 109
4.2.3 Choosing a galaxy parametrization . . . 113
4.2.4 The models . . . 118
4.2.5 Errors . . . 119
4.2.6 Timings . . . 120
4.3 Tests on Simulated Data . . . 120
4.3.1 Star formation histories . . . 121
4.3.2 Wavelength range . . . 123
4.3.3 Noise . . . 125
4.3.4 Dust . . . 127
4.4 Tests on real data . . . 129
4.4.1 Duplicate galaxies . . . 129
4.4.2 Real fits . . . 131
4.4.3 VESPA and MOPED . . . 132
4.5 Final remarks on VESPA . . . 137

5 A catalogue of star formation histories **141**

5.1 The Sloan Digital Sky Survey . . . 141

5.1.1 The main galaxy sample . . . 142

5.1.2 Spectro-photometric callibrations . . . 143

5.1.3 Galactic extinction . . . 143

5.1.4 Handling SDSS data . . . 143

5.2 The catalogue . . . 144

5.2.1 Masses and mass fractions . . . 145

5.2.2 Error estimates . . . 147

5.2.3 Mass and metallicity per age bin . . . 147

5.2.4 Dust . . . 148

5.2.5 V_{max} . . . 148

5.2.6 The database . . . 151

5.3 Basic results . . . 151

5.3.1 The observed galaxy stellar mass function . . . 153

5.3.2 The evolution of the GSMF . . . 155

5.3.3 Stellar mass density . . . 160

5.3.4 IMF studies . . . 163

6 Progenitors of supernovae type Ia **169**

6.1 Introduction . . . 169

6.1.1 SNIa progenitors . . . 171

6.1.2 SNIa evolution . . . 174

6.2 VESPA and SNIa . . . 175

6.2.1 Sample selection . . . 175

6.2.2 Method . . . 175

6.3 Results . . . 177

6.4 Interpretation . . . 177

6.4.1 Sample selection effects . . . 179

6.4.2 Comparing with previous results . . . 179

6.4.3 On the nature of the progenitors . . . 180

6.5 Future prospects . . . 181

7 Summary and future work **183**

7.1 Non-Gaussianity studies of the CMB . . . 183

7.2 VESPA . . . 185

7.3 Mass assembly in galaxies 186
7.4 Progenitors to SNIa 188
7.5 Final remarks 189

List of Figures

1.1 Hubble SNIa diagram from Perlmutter et al. (1999). 9
1.2 The matter power spectrum from a variety of sources (Tegmark et al. 2004) 19
1.3 The CMB temperature power spectrum, from Nolta et al. (2008) 28
1.4 Example of two different dust attenuation curves. 38

2.1 The contribution of the various foreground sources in the CMB. 49
2.2 Noise power-spectrum for different channels in the WMAP experiment. 52
2.3 The power-spectrum of a simulated Gaussian map, compared to the observed one. 55

3.1 Year one analysis: full-sky results. 63
3.2 Convergence tests for the full-sky analysis. 64
3.3 Year one analysis: North-South results. 64
3.4 $\Delta\xi_C(\theta)$ for the WMAP map high-pass filtered with different values of ℓ_{cut}. . . 66
3.5 Year one analysis: constraining in harmonic space. 67
3.6 Year one analysis: constraining in real and harmonic space. 68
3.7 Convergence tests for a selection of statistics. 68
3.8 Year one analysis: single-frequency analysis. 70
3.9 Year one analysis: removing the cosmological signal 72
3.10 Year five analysis: the auto-correlation function of peaks, on a full-sky analysis. 78
3.11 Convergence tests for the full-sky analysis. 78
3.12 Year five analysis: full-sky analysis results using the auto-correlation function. 79
3.13 Year five analysis: the cross-correlation function of peaks, on a full-sky analysis. 80

3.14 Year five analysis: full-sky analysis, using the cross-correlation function. 81
3.15 Year five analysis: convergence tests for a North-South analysis. 81
3.16 Year five analysis: auto-correlation function results, northern hemisphere. 82
3.17 Year five analysis: auto-correlation function results, southern hemisphere. 83
3.18 Year five analysis: cross-correlation function results, north-south analysis and constraints in real space. 84
3.19 Year five analysis: auto-correlation function results, constraining in real space - northern hemisphere. 85
3.20 Year five analysis: auto-correlation function results, constraining in real space - southern hemisphere. 86
3.21 χ^2 contributions towards signals in the V and W bands, per angular scale 87
3.22 The ISW contribution to the observed CMB, as calculated by Francis & Peacock 2008 (in prep). 88
3.23 The auto-correlation function for simulated maps with different values of f_{NL}. 94
3.24 The probability distribution of χ^2 for three values of f_{NL} assuming a true value of $f_{NL} = 40$. 95

4.1 The cosmic star formation history, from Hopkins and Beacom (2006). . 101
4.2 Evidence from downsizing in Panter et al. (2007) 102
4.3 The discrepancy between different estimations of ρ_* from Wilkins et al. (2008). 104
4.4 The behaviour of the singular values with matrix rank k. 113
4.5 The coefficients in sum (4.16) as a function of rank κ. 114
4.6 Schematic view of the grid of bins used by VESPA. 115
4.7 An example of an evolution of the fit, as VESPA searches for a solution. 117
4.8 Two examples of VESPA's analysis on synthetic galaxies. 120
4.9 The distribution of G_x, G_Z and total mass recovered for 50 galaxies with a SNR per pixel of 50. 121
4.10 Two examples of VESPA's analysis on synthetic galaxies, with reduced wavelength coverage. 122
4.11 The distribution of G_x, G_Z and total mass recovered for 50 galaxies with a SNR per pixel of 50 and two different wavelength coverages. 122
4.12 The number of recovered parameters changing with wavelength range. . 124

4.13 The distribution of G_x, G_Z and total mass recovered for 50 galaxies with different signal-to-noise ratios. . . . 125
4.14 The number of recovered parameters changing with SNR. . . . 126
4.15 Testing the SVD stopping criterion. . . . 127
4.16 Testing the recovery of τ_V^{ISM}. . . . 128
4.17 Average star formation history for two sets of duplicate galaxies in the SDSS. . . . 130
4.18 Stellar mass for two sets of duplicates galaxies in the SDSS. . . . 131
4.19 The distribution of reduced values of χ^2 for test sample of galaxies. . . . 133
4.20 Typical fit to a galaxy from the SDSS. . . . 134
4.21 Comparing the recovered star formation histories of a sample of galaxies using VESPA and MOPED. . . . 135
4.22 Comparing the stellar mass estimates for a sample of galaxies using VESPA and MOPED. . . . 136
4.23 Number of non-zero parameters in solutions recovered from a sample of galaxies with VESPA. . . . 137
4.24 The distribution of the total number of recovered stellar populations which contribute 5 per-cent or more to the total flux of the galaxy, as recovered from MOPED and VESPA. . . . 138

5.1 The measured galaxy stellar mass function of the SDSS's main galaxy sample for $0.005 < z < 0.35$. . . . 154
5.2 The measured galaxy stellar mass function of the SDSS's main galaxy sample for $0.005 < z < 0.35$ compared to other measurements from the literature. . . . 155
5.3 The SDSS DR5 galaxy stellar mass function split by redshift bins. . . . 157
5.4 The inferred evolution of the GSMF between $z = 0.9$ and $z = 0.5$ 159
5.5 The GSMF at $z = 0.5$ from VESPA (SDSS), GOODS-MUSIC, and COMBO-17 surveys. . . . 160
5.6 The observed evolution of the GSMF between $z = 0.9$ and $z = 0.5$ in the COMBO-17 data. . . . 161
5.7 Stellar mass density as a function of redshift. . . . 162
5.8 Dependence of $\rho_{*,SFH}$ and $\rho_{*,obs}$ on IMF parameters. . . . 164

5.9 The predicted local stellar mass density as derived from integrating the cosmic star formation history and from direct stellar mass measurements as a function of the IMF high-mass slope. 166

6.1 Prediction of SNR as a function of redshift, assuming an A+B model. . 173

6.2 SNIa rate per unit stellar mass, unnormalised, as a function of fraction of stellar mass. 178

List of Tables

1.1 Summary of the cosmological parameters derived from WMAP5 35

3.1 Year one analysis: a summary of non-Gaussianity detections. 74
3.2 The change in the value of $\chi^2_{reduced}$ for our hot spots statistics estimators due to the removal of the local ISW effect. 90
3.3 The change in the value of $\chi^2_{reduced}$ for our cold spots statistics estimators due to the removal of the local ISW effect. 91
3.4 The main indications of non-Gaussianity in the 5th-year data. 93

5.1 SDSS's fifth data release. 142
5.2 Galaxy properties which are derived by VESPA. 144
5.3 Additional properties derived from VESPA's output. 145
5.4 GalProp 151
5.5 BinProp 152
5.6 DustProp 152
5.7 BinID 152
5.8 Stellar mass density as a function of redshift. 162

6.1 Summary of A and B rates calculated in the literature. The two B rates from Scannapieco and Bildsten (2005) are calculated using two different estimators for the star formation rate. 172
6.2 Best-fit values of β/α for different aged populations and the respective significances. 178

Chapter 1

Introduction

Our current Standard Model is the Λ Cold Dark Matter model, in which the Universe went through a period of inflationary expansion at very early times. This expansion was driven by one or more scalar fields and inflated a small, causally-connected region of the Universe, to a size comparable to or larger than the size of the observable Universe today. The result was an almost uniform Universe, populated with small density fluctuations which grew under gravity to give rise to the structure we see today. We observe a hierarchical build-up of galaxies, with smaller galaxies merging together to form larger galaxies, of many and varied types. Furthermore, supernova observations have shown that the universe has recently become dominated by Dark Energy, causing its expansion to accelerate. This thesis probes the Universe at a variety of epochs, and aims to test observationally several stages of this model. In this Chapter we briefly summarise our current understanding of the Universe, from its content and large-scale dynamics, to the creation and evolution of structure and galaxies. We also describe two of the observables which are central to this thesis: the Cosmic Microwave Background and integrated galactic spectra.

1.1 The smooth Universe

The currently observed large-scale distribution of matter in the Universe has largely confirmed that we live in an isotropic and homogenous Universe, in accordance to the Cosmological Principle. Isotropy signifies that the Universe looks the same in all

directions, and homogeneity means that the Universe is the same everywhere. Treating the Universe as being uniformly smooth, homogeneous and isotropic allows us to gain insight on its dynamics as a whole, which in turn depend on its geometry and content. This section concerns the description of the Universe under these assumptions. We will see later that to explain the formation of structure that we observe today, we must allow for a departure from this assumption - this is discussed in section 1.4.

1.1.1 The geometry of the Universe

Homogeneity and isotropy, together with General Relativity, allow us to describe four-dimensional space time in the form of the Robertson-Walker metric

$$c^2 d\tau^2 = c^2 dt^2 - R(t)^2 \left[dr^2 + S_k^2(r) d\Psi^2 \right], \tag{1.1}$$

where $R(t)$ is the scale factor and has units of distance, r is a dimensionless comoving distance, k is the curvature constant, Ψ is the angular separation between the two events, τ is the proper time and t is the cosmological time. The form of $S_k(r)$ depends on the curvature

$$S_k(r) = \begin{cases} \sin(r) & \text{if } k = 1 \\ \sinh(r) & \text{if } k = -1 \\ r & \text{if } k = 0. \end{cases}$$

The curvature constant describes the local geometry of the Universe: $k = 1$ implies a closed Universe, $k = -1$ an open Universe, and $k = 0$ a flat Universe.

1.1.2 The dynamic Universe

Expansion

In 1929 Edward Hubble observed a tenuous but true correlation between the distance of a galaxy, d, and its apparent recession velocity from the Earth, v

$$v = Hd. \tag{1.2}$$

This led him to infer that the Universe is expanding, going against the view at the time that we lived in a static Universe. The rate of this expansion is nowadays represented by the Hubble parameter $H(t)$ and

$$H(t) = \frac{\dot{R(t)}}{R(t)}. \tag{1.3}$$

Here the dot is used to represent the derivative with respect to cosmological time, t - this notation is used throughout this thesis. Constraints on the current value of the rate of expansion put $H_0 = 71.9$ km s^{-1} Mpc $^{-1}$ (Komatsu et al. 2008).

Redshift

An immediate consequence of an expanding universe is that we expect a change in the frequency of a light signal caused by the relative velocity between the light source and the observer. We define the redshift, z as the change in frequency of the emitted and observed signal:

$$\frac{\nu_{em}}{\nu_{obs}} = 1 + z \tag{1.4}$$

For small changes in frequency we can write this in terms of the Doppler shift, which relates the change in frequency with the velocity of the source. This in turn can be associated with the Hubble parameter and the scale factor:

$$\frac{\delta\nu}{\nu} = \frac{\delta v}{c} = -\frac{Hd}{c} = H\delta t = -\frac{\dot{R}}{R}\delta t = -\frac{\delta R}{R}. \tag{1.5}$$

The minus sign in the second step arises from the fact that the observer and the source have opposite relative velocities (i.e., they are receding from one another) and gives rise to the result

$$\nu \propto \frac{1}{R}. \tag{1.6}$$

Practically we are concerned with the shift in frequency for a signal emitted by an object at a redshift z and measured by us, sitting at $z = 0$. Combining equations (1.4) and (1.6) we can write

$$\frac{R_0}{R(t)} = 1 + z. \tag{1.7}$$

For convenience we also define a normalised scale factor as $a(t) = R(t)/R_0$.

The Friedmann equation

The dynamics of the Universe can be described by the Friedmann equation, which relates the energy content of the Universe with its dynamical evolution and geometry and arises from solving Einstein's equations:

$$\dot{R}^2(t) - \frac{8\pi G}{3}\rho(t)R^2(t) = -kc^2, \tag{1.8}$$

where G is Newton's constant of gravity. We can define a critical density, ρ_{crit}, which marks the density needed for a flat Universe and therefore the transition from the closed to the open case. By setting $k = 0$ in equation (1.8) and using equation (1.3)

$$\rho_{crit}(t) = \frac{3H(t)^2}{8\pi G}. \tag{1.9}$$

It is common to express the energy content of the Universe as a ratio in relation to ρ_{crit}:

$$\Omega(t) = \frac{\rho(t)}{\rho_{crit}(t)}. \tag{1.10}$$

The present-day value of the energy density, scale factor and the Hubble parameter are normally denoted with a 0 subscript, e.g. H_0. In this thesis we will drop the subscript from Ω_0, and explicitly write-down the time or redshift dependence when we are referring to its value at some time other than the present-day.

Taking the derivative of equation (1.8) with respect to time gives an expression for the acceleration of the Universe. To do this we need to know how the density evolves with time, but using an argument of conservation of energy coupled with the assumption that the expansion of the Universe is adiabatic we can write $dE = -pdV$, where E is the total energy, p is pressure, and V is the expanding volume $V \propto R^3$. We can write

$$\ddot{R} = -\frac{4\pi}{3} GR \left(\rho + \frac{3p}{c^2} \right). \tag{1.11}$$

Formally, both equation (1.8) and (1.11) arise independently from a full treatment using General Relativity and the Robertson-Walker metric. Written in this form, we see how the pressure can act as an extra form of gravity which is not an intuitive result. For the moment we simply note that the contribution of pressure is important in many cosmological applications, as we will see later.

1.1.3 The equation of state

Solving Friedmann's equation gives the evolution of the expansion of the Universe at large-scales. To do this, we need an explicit form for the energy density, and that in turn depends on its nature. Pressureless matter has an energy density which goes as $\rho_m \propto R^{-3}$, because number density of particles must be conserved within a given comoving volume. Radiation's energy density contribution loses an extra power of R, arising from the redshifting of energy ($E = h\nu$), and we have $\rho_r \propto R^{-4}$. The steeper

dependence of ρ_r on the expansion, compared to ρ_m tells us immediately that, for a small enough R, radiation dominated over matter. We will review this epoch of the Universe in section 1.5.

We also consider the existence of a vacuum energy component. First introduced by Einstein in order to explain the static Universe which was at the time observed, it was then later dropped when improved observations revealed a Universe which was expanding. Today, a non-negative vacuum energy density, ρ_v is invoked to explain the fact that the expansion of the Universe is accelerating (see section 1.2.4 and Figure 1.1). For the moment we need only consider that, as a property of empty space, ρ_v is constant and has no dependence on R.

We can now write a general form for Friedmann's equation, which includes the contributions from the three components mentioned above. Using $\rho = \rho_{0,m}\left(\frac{R}{R_0}\right)^{-3} + \rho_{0,r}\left(\frac{R}{R_0}\right)^{-4} + \rho_v$, together with equations (1.8) and (1.3) and in terms of critical densities as defined in (1.10) we write

$$H^2(z) = H_0^2\left[\Omega_v + \Omega_m(1+z)^3 + \Omega_r(1+z)^4 - (\Omega - 1)(1+z)^2\right], \tag{1.12}$$

where the final term comes directly from the curvature term in Friedmann's equation.

1.2 The components of the Universe

The Universe is composed of a mixture of radiation, matter and Dark Energy. As we have seen, each of these components has a different dependence on the scale factor, R. What this means is that the history of the Universe is dominated by different components at different times, resulting in dramatic changes throughout its lifetime.

1.2.1 Radiation

As we have seen in section 1.1.3, we expect that at very early times the Universe was dominated by radiation. If we assume an adiabatic expansion (by which we mean the entropy change in any comoving region is zero), then $T \propto V^{(1-\gamma)}$, where γ is the ratio of specific heats and is equal to 4/3 for radiation. Therefore we get $T \propto R^{3(1-\gamma)}$, which simply gives $T \propto 1/R$ - i.e. the Universe was very hot at very early times.

The temperature and rate of expansion in the early Universe dictate the abundance of photons (γ) and neutrinos (ν) we see today. Briefly, we generally assume that at early times the Universe was in thermal equilibrium. As the expansion develops the temperature drops and, one by one, the reactions that keep each of the species in equilibrium cease when the interaction time scale is longer than the expansion rate. When the temperature reaches 10^{10} K, the only species still in equilibrium are photons, neutrinos and electron-positron pairs. The latter annihilate when the temperature reaches $10^{9.7}$ K, in a reaction which creates an excess of photons relative to neutrinos. Conservation of entropy requires (Peacock 1999):

$$T_\nu = \left(\frac{4}{11}\right)^{1/3} T_\gamma. \tag{1.13}$$

We will see later that today we measure $T_\gamma = 2.725$ K in the Cosmic Microwave Background - the photon radiation relic from the Big Bang. This implies a neutrino background with $T_\nu = 1.94$ K.

If neutrinos and photons are the only contributions to the radiation energy density then the redshift for the matter-radiation equality, z_{eq}, is (Peacock 1999):

$$1 + z_{eq} = 23900\Omega h^2 \left(\frac{T}{2.73K}\right)^4. \tag{1.14}$$

We will see later that this is a very important epoch for the formation of structure in the Universe. Its value is observationally constrained to be approximately 3100. At this point the temperature was still high enough that atoms were fully ionized and matter was coupled to radiation via Thomson scattering. Matter and radiation finally decouple at a redshift $z_{dec} \approx 1100$. This marks another crucial moment in the evolution of the Universe, and the creation of the Cosmic Microwave Background.

1.2.2 Baryonic matter

The high temperature in the early Universe, which up to a point exceeded that found at the centre of stars, suggests that primordial nucleosynthesis must have happened. Whilst the temperature is high enough, protons and neutrons are in thermal equilibrium. Once the temperature is low enough, they combine to form nuclei - this happens when the temperature reaches $\approx 10^{10}$ K. The first element to form is Deuterium, which can in turn combine to form Helium. The temperature drops too quickly before any significant amount of other nuclei has the opportunity to form. The wide variety of

elements we are familiar with today, comes from nuclear reactions at the centre of stars or supernovae explosions. Observed Deuterium abundances can be used to estimate the baryon critical density Ω_b, since Deuterium is normally destroyed in stellar nuclear reactions. Estimates for Ω_b from primordial nucleosynthesis and other methods (e.g. CMB) are in good agreement and give $\Omega_b \approx 0.04$.

1.2.3 Dark Matter

There is solid evidence for a matter component beyond baryonic matter. Some form of unseen matter which would respond to gravity was proposed first by Zwicky (1933), to explain galaxy velocities in the Coma cluster. Further evidence since then includes the rotation curves of galaxies, which are much flatter than the $1/r^2$ expected from the luminous component, and the recently found Bullet Cluster (Clowe et al. 2006), which clearly shows a separation between the baryonic gas and the dark component. Despite our ignorance about its nature, Dark Matter is now an important part of our current model of galaxy formation which, as we will see, we also do not yet completely understand. Advancement in our understanding on the nature of Dark Matter is most likely to come from underground experiments which aim to detect Dark Matter particles directly (e.g. the Boulby mine project in North Yorkshire, Paling 2005).

Even though they are presently not favoured by observations, a lot of work is going into developing modified theories of gravity. First proposed by Milgrom (1983), modified theories of gravity propose a change to the laws of gravity, either in the framework of Newtonian dynamics or General Relativity, in an attempt to explain the Universe without the need for Dark Matter. Very many theories have been proposed since 1983 but none, so far, explains the observable Universe as well as the assumption of Dark Matter.

1.2.4 Dark Energy

The first evidence that our Universe is dominated by Dark Energy came from the observation of distant type Ia supernovae (SNIa). The experiment is conceptually very simple: if we can find a standard candle in the Universe, then by measuring its apparent brightness we can calculate its distance to us. As we will see, this will depend both on the redshift and but also on the cosmological model and this allows us to constrain the geometry of space between us and the standard candle. SNIa are very good candidates for standard candles, and in 1998 and 1999 two experiments revealed that distant

SNIa are much dimmer than was expected at the time (Perlmutter et al. 1999; Riess et al. 1998), implying these objects are at a greater distance from us. Figure 1.1, from Perlmutter et al. (1999), shows how the data is able to differentiate between different cosmological models for intermediate to high redshift. In practice, SNIa are not perfect standard candles, and demand a great deal of care when making plots such as the one in Figure 1.1 - see Chapter 6 for more details.

The SNIa experiments require the Universe's expansion to have become *accelerated* in recent times. Recall equation (1.11). We see immediately that for $\ddot{R} > 0$ we require $\rho + 3p/c^2 < 0$. We therefore require a form of energy density which has negative pressure. This energy component of the Universe has been named Dark Energy because we do not know what it is. One of the goals of modern cosmology is to constrain its equation of state

$$p_v = w\rho_v c^2. \tag{1.15}$$

$w = -1$ would correspond to Einstein's cosmological constant, with which observational constraints are consistent. From the condition $\rho + 3p/c^2 < 0$, we immediately get an upper limit of $w < -1/3$. Assuming a flat Universe and a constant equation of state with redshift CMB constraints yield $w = -0.967^{+0.073}_{-0.07}$ (Spergel et al. 2007), SNIa experiments $w = -1.07^{+0.09}_{-0.09}$ (Wood-Vasey et al. 2007), and a combination of large-scale structure with SNIa and CMB gives $w = -1.004 \pm 0.089$ (Percival et al. 2007a), to name only a few.

1.3 Observational cosmology tools

It is useful to summarise some of the relations we derived, in the context of observational cosmology. As observers, we measure the redshift z and angular sizes or distances in the sky, $d\psi$, and we would like to relate them to quantities such as size or distance, volume and age. Let us start by relating time and redshift. By taking the derivative of equation (1.7) with respect to time we get

$$dz = -(1+z)H(z)dt, \tag{1.16}$$

where $H(z)$ is defined in equation (1.12). Integrating this relation between the appropriate limits then gives us either the age of the Universe at a redshift z, or the

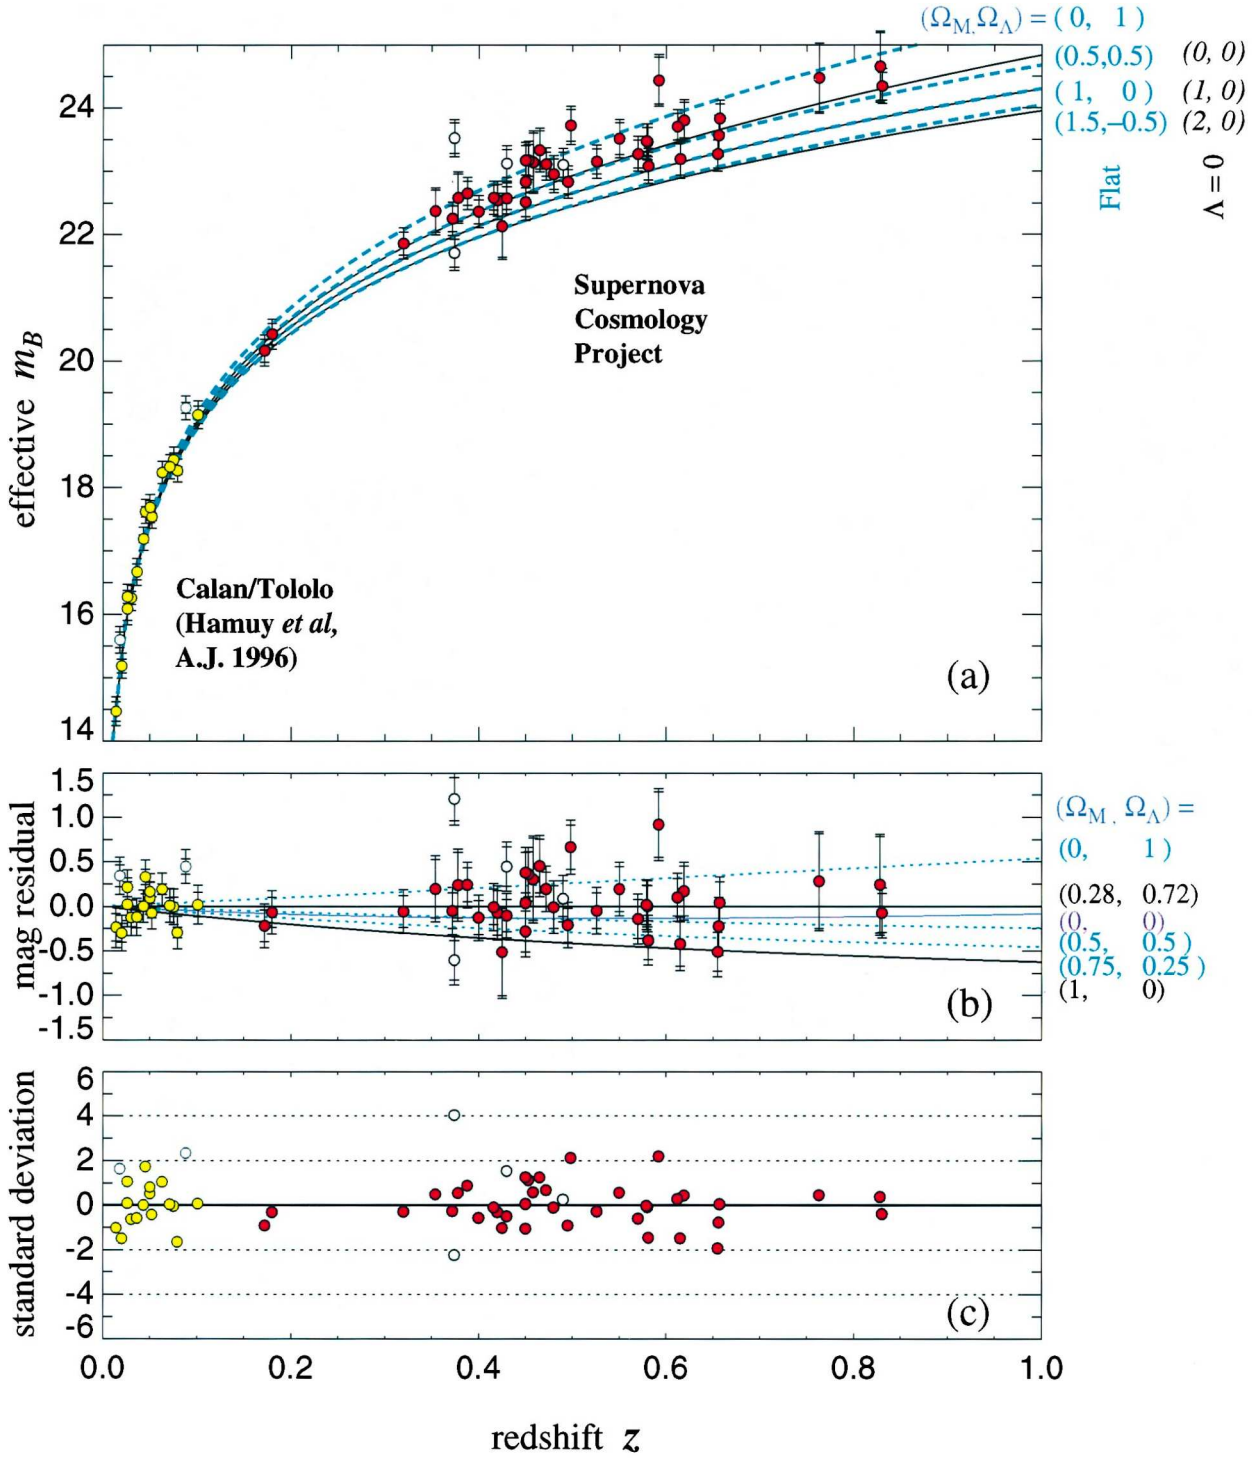

Figure 1.1: From Perlmutter et al. (1999): panel a) shows evidence for an accelerated expansion using SNIa; panels b) and c) show residuals for the cosmological fit.

lookback time since redshift z.

To get the relation between redshift and comoving distance, we first need the equation of motion of a photon, which relates r and t. General Relativity tells us that photons move along geodesic paths ($d\psi = 0$) and have zero proper time $d\tau = 0$. It can be seen immediately from the Robertson-Walker metric (equation 1.1) that $cdt = R(t)dr$. Using equation (1.16) gives:

$$R_0 dr = \frac{c}{H(z)} dz. \tag{1.17}$$

This is a particularly useful relation in observational cosmology, as it allows us to calculate sizes and volumes. Consider the Robertson-Walker metric once more. The spatial part can clearly be divided into a radial and an angular component, both of which are parametrized by the scale factor to account for the expansion. The proper transverse size of an object is given by its angular component:

$$d\ell_\perp = d\psi R(z) S_k(r) = d\psi R_0 S_k(r)(1+z)^{-1} \tag{1.18}$$

We are interested in the scale factor at the redshift of the observation because this is the proper size of an object - and that is not affected by the expansion. We can combine the radial and angular parts to write down the volume element of a shell of area $d\psi^2$ and comoving radius dr to get:

$$dV = R^3(z) S_k^2(r) d\psi^2 dr. \tag{1.19}$$

In this case we are generally interested in the comoving volume, since this is the volume in which number densities of galaxies remain constant in the Hubble flow:

$$dV_c = R_0^3 S_k^2(r) d\psi^2 dr^2. \tag{1.20}$$

We now define two measurements of distance as a function of redshift, each defined as an attempt to connect our notions of Euclidean space with a RW space. One definition comes directly from equation (1.18): we can see by inspection that it is very similar to what we would expect from Euclidean geometry (i.e. $d\ell_\perp = D_A d\psi$) if we define

$$D_A = \frac{R_0 S_k(r)}{1+z}. \tag{1.21}$$

The other definition comes from considering the observed bolometric flux, F_{tot} of a source at a redshift z, with an assumed power law spectrum $L \propto \nu^{\alpha}$ and total luminosity L_{tot} (more details in Peacock (1999) and section 5.2.1):

$$F_{tot} = \frac{L_{tot}}{4\pi R_0^2 S_k^2(r)(1+z)^2}. \tag{1.22}$$

Once again we make the observation that equation (1.22) can a take a similar form for that obtained in Euclidean space ($F_{tot} = L_{tot}/(4\pi D_L^2)$) if we define:

$$D_L = R_0 S_k(r)(1+z) \tag{1.23}$$

All of the relations in this section are model dependent. This is not a problem in practice - nowadays cosmological parameters are well constrained (see Table 1.1), but this demands care when comparing observational results across different cosmologies.

1.4 Structure formation

The Universe described in section 1.1 is a smooth Universe, and although it provides a good description of the real Universe on large scales, it crucially fails to explain the large gradients in density we see today. In this section we will see how structure is seeded and how it evolves. A full treatment is a technical challenge, and here we concentrate on important results which give an insight on how different physical mechanisms shape the evolution of the Universe as it goes through key stages of this process. Firstly we will introduce Inflation as a mechanism which introduces small density fluctuations in the Universe. If these fluctuations are small, one can use linear perturbation theory to follow their growth until the time they enter the non-linear regime. After this stage there are no analytical models to describe the evolution of the perturbations, and we mostly rely on numerical simulations for accurate answers. We will see however, that there are analytical approximations which provide recipes to treat perturbations in the non-linear regime and which give an insight on how galaxies form and on how matter is distributed in the Universe.

1.4.1 Inflation

Inflation was introduced in 1981 by Alan Guth (Guth 1981) as an early Universe theory which aimed to solve what was known as the horizon problem. This arises because regions of sky which today are further apart than roughly one degree in the sky were not

within causal contact at the time the Cosmic Microwave Background was created (see section 1.5.2 for more details), therefore providing no apparent reason as to why the Universe seems to be homogeneous on larger scales. Inflation solves this problem by proposing an accelerated expansion at early times, which quickly and briefly expanded a small region of Universe to a size at least as large as the size of the observable Universe today. The attraction of Inflation is that it also solves three more problems with the standard Big Bang model.

Firstly, it solves the flatness problem: the fact that we observe a Universe very close to flat requires it to have been very close to flat in the past, which suggests some fine-tuning. Consider Friedmann's equation in the case of a vacuum-dominated Universe:

$$\dot{R}^2 = \frac{8\pi G\rho_v R^2}{3} - kc^2. \tag{1.24}$$

As ρ_v has no dependence on R, we have a simple solution for R as

$$R \propto \exp\left(\pm\sqrt{\frac{8\pi G\rho}{3}}\right). \tag{1.25}$$

The exponential expansion with time means that the term $\rho_v R^2$ will dominate over the curvature term until it becomes negligible - making the Universe tend to flat. In this argument we also connected the idea of Inflation with the idea of vacuum energy density. The connection arises because Inflation requires accelerated expansion which is precisely the behaviour a vacuum energy term gives.

In doing so, we are solving another problem with the classic Big Bang theory - Inflation provides us with a *reason* as to why the Universe is expanding today. Even though the Inflationary period had to be brief, it gave the Universe an initial momentum which, in the absence of a vacuum energy dominated era in recent times (see Section 1.2.4), is enough to explain its current expansion.

Finally, Inflation takes us from a smooth Universe into one populated by density fluctuations. To understand how, we need to briefly look at the nature of the physics behind Inflation. Most commonly, Inflation is associated with a scalar field, ϕ which is quantum in nature, and the associated potential $V(\phi)$. For a homogenous field we can write (Dodelson 2003)

$$\ddot{\phi} + 3H\dot{\phi} + \frac{\partial V}{\partial \phi} = 0 \tag{1.26}$$

with

$$\rho = \frac{1}{2}\dot{\phi}^2 + V(\phi) \quad \text{and} \quad p = \frac{1}{2}\dot{\phi}^2 - V(\phi). \tag{1.27}$$

To have an inflationary behaviour, we want a field which has $\rho + 3p < 0$. This condition is normally cast in the following relations (Peacock 1999):

$$3H\dot{\phi} = \frac{-\partial V}{\partial \phi} \tag{1.28}$$

$$\epsilon \equiv \frac{m_P^2}{16\pi}\left(\frac{V'}{V}\right)^2 \ll 1 \tag{1.29}$$

$$\eta \equiv \frac{m_P^2}{8\pi}\left(\frac{V''}{V}\right) \ll 1 \tag{1.30}$$

We are mostly interested in the output of this framework, in terms of its observables. The first thing to appreciate is that statistical quantum fluctuations in the inflation field, $\delta\phi$ create scalar perturbations in the metric, which ultimately give rise to the inhomogeneities in the gravitational potential, $\delta\Phi$, needed to seed cosmic structure. These are the perturbations we are mostly interested in for this thesis. However, tensor fluctuations in the gravitational metric are also expected, and these in turn give rise to a background of gravitational waves. This background is yet undetected, and if detected would provide one of the best pieces of evidence for Inflation. Dodelson (2003) provides a clear treatment for both cases. The mean of the scalar perturbations is zero at any given time, but we are mostly interested in its variance $\langle\Phi^2\rangle$ - essentially the power-spectrum of the resulting fluctuations after Inflation has come to an end. These fluctuations are expected to be Gaussian distributed around zero, which is an important aspect we will return to in more detail.

We define the power-spectrum of perturbations as

$$\langle\delta_{\mathbf{k}}\delta^*_{\mathbf{k}'}\rangle \equiv (2\pi)^3 P(k)\delta(\mathbf{k} - \mathbf{k}'). \tag{1.31}$$

and we write the primordial power-spectrum as (Dodelson 2003)

$$P_\Phi(k) = \frac{50\pi^2}{9k^3}\left(\frac{k}{H_0}\right)^{n-1}\delta_H^2\left(\frac{\Omega_m}{D_{1(a=1)}}\right) \tag{1.32}$$

where δ_H defines the scalar perturbation amplitude at the time of horizon crossing, and D is the growth function (see next section). n is called the spectral index and is related to the potential $V(\phi)$ by the quantities η and ϵ defined in equations (1.29) and (1.30). In practice, most theoretical models predict n to be close to, but not necessarily one. We will see in the next section how we cannot observe this primordial spectrum directly because its shape is changed by the evolution of the density fluctuations after Inflation has ended. However, if we understand the growth of these fluctuations with time, we can understand the observed matter power-spectrum as to allow us to constrain the value of n and give insight on the shape of the potentials which are still unconstrained from a theoretical (and observational) point of view. We will refer to a spectrum in which $k^3 P_\Phi(k) =$ constant ($n = 1$) as a *scale-invariant spectrum*, or a Harrison-Zel'dovich spectrum, after the two people who first suggested a spectrum of this form (long before the idea of Inflation).

We will see that these fluctuations will also propagate themselves into the primordial radiation field, which we can observe today as the Cosmic Microwave Background. Section 1.5 looks at the origin of the CMB in more detail, but for now we would like to stress the point that the propagation of $\delta\Phi$ to the observed temperature fluctuations on the CMB is predicted to be linear or nearly-linear, which in turn means that temperature fluctuations will also exhibit Gaussian statistics.

Perturbations in the density field are related to perturbations in the gravitational potential by Poisson's equation: $\nabla^2 \delta\Phi = 4\pi G \delta\rho$. Consider the two main components of the energy density: matter and radiation. There are two types of perturbations which are normally considered, which relate these two quantities in different ways. For an *adiabatic* perturbation, with $T \propto 1/R$ and constant entropy, we have $\delta_r = \frac{4}{3}\delta_m$. For an isocurvature perturbation, the total change in the energy density is zero and $\delta_r \rho_r = -\delta_m \rho_m$ (Peacock 1999). Observationally, adiabatic perturbations are favoured to isocurvature perturbations (e.g. Efstathiou and Bond 1986).

1.4.2 The linear regime

Inflation left us with perturbations in the density field around a background which is smoothly and uniformly expanding. In this section we will delineate the formal treatment generally used to follow the evolution of these perturbations, which combines linear perturbation theory and fluid dynamics. Particularly clear derivations, which fill

the gaps between the key steps we mention next can be found in Peacock (1999) and Binney and Tremaine (2008).

We will be working with density fluctuations, defined as:

$$1 + \delta(\mathbf{x}) \equiv \frac{\rho(\mathbf{x})}{\langle\rho\rangle} \tag{1.33}$$

and we will mainly work in Fourier space, as then the modes grow independently:

$$\delta(\mathbf{x}) = \frac{1}{V}\sum_k \delta_\mathbf{k} e^{i\mathbf{kx}} \quad \text{and} \quad \delta_\mathbf{k} = \int_V d^3\mathbf{x}\delta(\mathbf{x})e^{-i\mathbf{kx}}. \tag{1.34}$$

The fluid dynamics equations behind this treatment are: a) the Continuity equation which tells us that the total mass must be conserved; b) the Euler equation, which tells us what the acceleration due to pressure gradients and gravity is; and c) Poisson's equation. For the case of collisionless dark matter we reach the following relation:

$$\ddot{\delta}_\mathbf{k} + 2H\dot{\delta_\mathbf{k}} = 4\pi G\rho_0\delta_\mathbf{k}. \tag{1.35}$$

If we for a moment ignore the term $2H\dot{\delta_\mathbf{k}}$, the equation has a simple exponential solution, as $\exp(\pm\sqrt{4\pi G\rho_0})$. Formally, every linear combination of the two solutions is also a solution to the initial differential equation. However, given the context we are interested in the *growing modes* so we will concentrate on these. What we find is that the perturbations grow exponentially under gravity. The effect of re-introducing the expansion term $2H\dot{\delta_\mathbf{k}}$, is to slow down this collapse and the solutions are now more like power-laws. For this reason this is normally called the *damping factor*. For a matter-dominated Universe, with $\Omega_m = 1$, the growing mode is

$$\delta(t) \propto t^{2/3} \propto a(t). \tag{1.36}$$

If we now add photons to our fluid, we are effectively introducing a pressure term, via the Euler equation, which is related to the density through the speed of sound $c_s = \partial p/\partial\rho$. This gives:

$$\ddot{\delta}_\mathbf{k} + 2H\dot{\delta}_\mathbf{k} = \dot{\delta}_\mathbf{k}\left(4\pi G\rho_0 - c_s^2\frac{k^2}{a^2}\right). \tag{1.37}$$

We can identify two regimes, in which either gravity or pressure dominate. The scale at which the two terms are balanced is called the *Jeans length*, $\lambda_J = c_s\sqrt{\frac{\pi}{G\rho}}$ which, for the radiation-dominated era, is of the order of the horizon size. The growth of the perturbations is now qualitatively very different for small and large scales. Whereas in

a $\Omega_m = 1$ Universe $\delta(t)$ continues to grow as $\delta(t) \propto t^{2/3}$ for scales larger or comparable to the size of the horizon, on smaller scales pressure acts as a restoring force which stops the collapse and sets up oscillations in the fluid. The general solution can be written as $\delta(a) \propto af\left(\Omega(a)\right)$ with $f(\Omega)$ approximated as (Carroll et al. 1992):

$$f(\Omega) \simeq \frac{5}{2}\Omega_m \left[\Omega_m^{4/7} - \Omega_v + (1 + \frac{1}{2}\Omega_m)(1 + \frac{1}{70}\Omega_v)\right]^{-1}. \tag{1.38}$$

At early enough times, we have seen that the Universe was dominated by radiation and the analyses we have used so far are no longer valid. The reason is that for a radiation fluid the mass-density continuity equation no longer applies: the total energy of a body of radiation *decreases* with expansion. A full relativistic treatment is needed, or a short-cut using the conservation of entropy can be found in Binney and Tremaine (2008). Here we will simply write down the equation of motion:

$$\ddot{\delta}_{\mathbf{k}} + 2H\dot{\delta}_{\mathbf{k}} = \delta_{\mathbf{k}}\left(\frac{k^2c^2}{3a^2} - \frac{32\pi}{3}G\rho_0\right). \tag{1.39}$$

Similarly to what happened in the baryon-fluid case, we find that for scales smaller than the horizon we expect radiation pressure and gravity to set up oscillations in the fluid: sound waves. For large scales:

$$\delta(t) \propto t. \tag{1.40}$$

The baryonic component of the energy density has, at this stage, little influence on the evolution of the perturbations and δ_m follows δ_r. Baryonic matter is fully ionized, and is coupled to the radiation through Thompson scattering (which couples the photons to the electrons) and Coulomb interactions (which couples the electrons to the baryons). Perturbations in collisionless dark matter are also prevented from collapsing in sub-horizon scales, but clearly the reason must be something other than radiation pressure. In this case this happens because the expansion rate is faster than the characteristic growth time for dark matter, and fluctuations freeze.

Let us summarise this section by identifying three key stages in the evolution of the perturbations:

- *Radiation-dominated era:* matter and radiation are coupled through Thomson scattering. $\delta(t) \propto t$ on scales larger than Jeans length, which at this stage is of the order of the horizon size. On scales smaller than the Jeans length radiation and matter are prevented from collapsing further due to radiation pressure - this

sets up oscillations. Dark matter perturbations are frozen on these scales due to the rate of expansion.

- *Radiation-matter equality:* at $z_{eq} \approx 3100$, the Universe becomes matter-dominated. $\delta \propto t^{2/3}$ both for matter and radiation (which are still coupled) and now also for dark matter, given that the rate of expansion slows down.

- *Decoupling:* at $z_{cmb} \approx 1100$ radiation and matter de-couple and evolve separately. Photons are no longer trapped, and free stream. Dark matter continues to self-gravitate, and baryonic matter traces dark matter from now on, due to gravity.

The description we gave above is a simplification, even in the context of linear theory. All contributions to the energy density are coupled and do not form a simple fluid. The overall physics which shapes the evolution of the perturbations is normally encapsulated in the transfer function, which we define as

$$T_k = \frac{\delta_{\mathbf{k}}(z=0)}{\delta_{\mathbf{k}}(z)D(z)} \tag{1.41}$$

where $D(z)$ is called the growth factor which traces the linear evolution of the perturbations.

1.4.3 The matter power-spectrum

Let us define a useful dimensionless form for the power-spectrum as

$$\Delta^2(k) = \frac{k^3 P(k)}{2\pi^2}. \tag{1.42}$$

The matter power-spectrum we measure today is the result of a matter density distribution described by a Harrison-Zel'dovich spectrum at the end of Inflation and which is subsequently changed by effects such as gravitational collapse and pressure. We introduced the transfer function in equation (1.41) as a short-hand to write these effects. The observed power-spectrum is therefore

$$\Delta^2(k) \propto k^{3+n} T_k^2 \tag{1.43}$$

Mainly due to practical reasons, we are often interested in the density field convolved with a Gaussian or a top-hat spherical function, $W_k(R_s)$ with an associated radius R_s. The crucial quantity here is the rms of this quantity, given by

$$\sigma^2(R_s) = \int \Delta^2(k)|W_k(R_s)|^2 d\ln k. \tag{1.44}$$

This also provides a route to an empirical normalization of the linear, matter power-spectrum, which is unconstrained by theory, through its value for $R_s = 8$ Mpc/h. Constraints from WMAP5 put this value at $\sigma_8 = 0.796 \pm 0.036$ (Komatsu et al. 2008).

Figure 1.2 shows the observed matter power-spectrum, estimated from a variety of sources each probing different scales. The turn in the power-spectrum corresponds to the horizon size at the time of matter-radiation equality. As we saw in the previous section, at this epoch perturbations at sub-horizon scales see their growth damped by pressure terms. As the Universe expands, different scales enter the horizon. The smallest scales enter the horizon first, and therefore have their growth more damped in relation to the other scales, which continue to grow for longer. An imprint of these acoustic oscillations can also been seen in the galaxy distribution, although the signal is mostly hidden by the data points in Figure 1.2. This signal has now been detected both in the Sloan Sky Digital Survey (SDSS) (e.g. Percival et al. 2007b) and in the 2dF Galaxy Survey (Cole et al. 2005).

1.4.4 The hierarchical model

Effectively, the transfer function acts to reduce the amplitude of the small scale perturbations via two main mechanisms: Jeans-mass effects as we have seen above, but also through *damping*. At very early times dark matter particles will be highly relativistic and free stream without much trouble, erasing any fluctuations on scales below the horizon size at that time. The time at which this ceases to happen is of crucial importance for the nature of structure formation. For massive particles, such as *cold dark matter*, this will happen long before the matter-radiation equality time, and scale fluctuations smaller than the horizon size at z_{eq} are able to survive. For *hot dark matter*, such as massive neutrinos, this only happens at z_{eq}. The result is that only fluctuations larger than the size of the horizon at z_{eq} are able to survive. To explain the fact that today's observed power-spectrum sees fluctuations below that scale, one must invoke a *top-bottom* scenario: i.e., galaxies formed from the dissipation of larger structures. The cold dark matter scenario however, appeals to a *bottom-up* growth, with the smaller scales being the first to collapse after z_{eq} and then merging to form larger structures, in what we call a *hierarchical* model. This can also be seen from equation (1.44). For a top-hat spherical filter, and $n = 1$ we find $\sigma^2(R) \propto R^{-2.5}$ - i.e., smaller scales are the first to collapse.

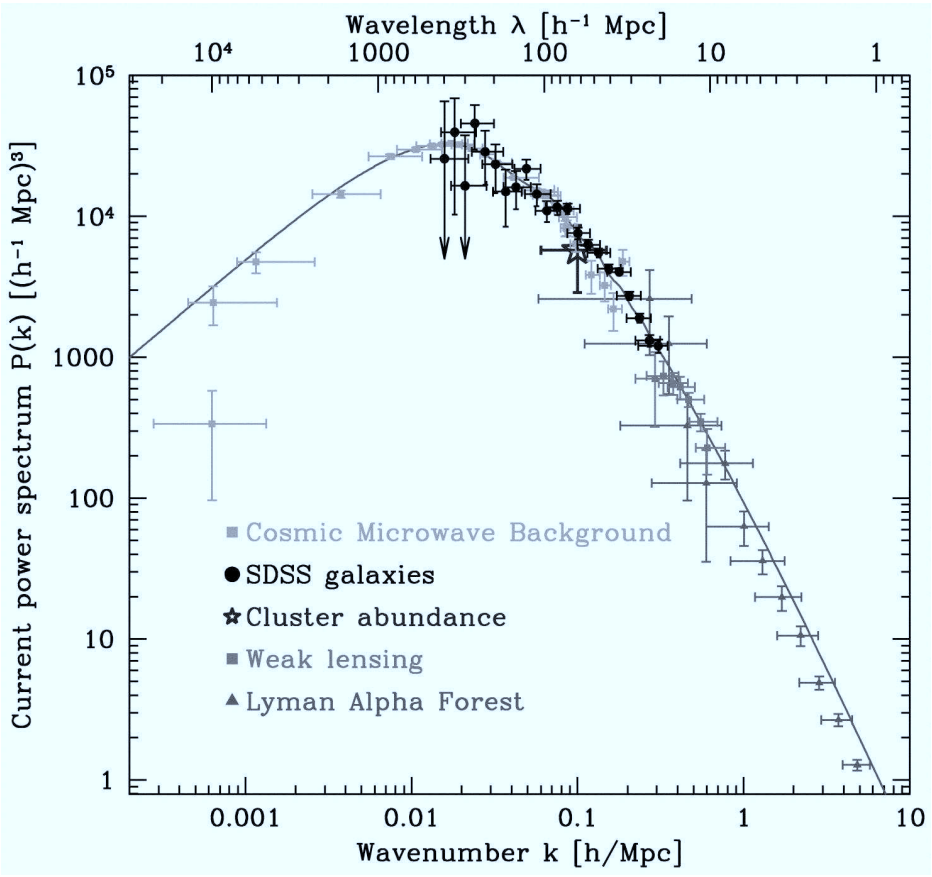

Figure 1.2: From Tegmark et al. (2004): the observed matter power-spectrum, measured at a variety of scales using different physical probes.

1.4.5 Beyond the linear regime

When the density perturbations become too large, linear perturbation theory is no longer valid. The density in a luminous red galaxy is roughly 10^5 times larger than the critical density ρ_{crit} (equation 1.9), so to explain the growth of structure to that level we are beyond the realm of linear perturbation theory. An exact answer of the growth of perturbations up to high densities can only be achieved with numerical simulations. Here we will briefly outline one of the many analytical approximations to the problem.

We are interested in tracking the growth of the perturbations such that we can predict real observables, which we can choose to look for. Perhaps two of the most fundamental properties about a galaxy are its luminosity and its mass. Let us then start then by deriving a mass function $n(M)$, defined such that $n(M)dM$ is the comoving number density of objects of mass in $[M, M + dM]$. We will also discuss the implications in terms of the luminosity function which is defined in a similar way, i.e., the comoving number density of galaxies with luminosity in $[L, L + dL]$.

To do this we will need to use results from two standard formalisms: the spherical collapse model, and the Press-Schechter theory (Press and Schechter 1974). Detailed descriptions of these can be found in Peacock (1999); Coles and Lucchin (1995) and of course, Press and Schechter (1974). Here we will simply quote the results that we need.

The spherical collapse model tracks the growth of a perfectly spherical perturbation with constant density inside of it. This spherical inhomogeneity sits in a smoothly expanding background, and we aim to track its evolution with time. The symmetry of the situation means this perturbation can be treated as an isolated closed Universe and Friedmann's equations apply (e.g. Peacock (1999)). Its evolution has three distinct phases: it initially expands with the background, then stops, collapses and finally virializes at a time t_v. What we want to know is the value of δ, in the linear regime, at the time t_v. We will call this value the critical overdensity for collapse, δ_c. The density of the perturbation in the non-linear regime will be greater that δ_c, but we are interested in *identifying* the regions in the density field which should undergo gravitational collapse. We can do this by seeing when $\delta(\mathbf{x}) = \delta_c$.

For each mass M, we can identify a scale R_s which corresponds to the radius of a sphere which contains a mass M, assuming a uniform background mean density ρ_0,

$M = \frac{4\pi}{3}\rho_0 R^3$). We have already seen that δ is Gaussian distributed with the rms given by equation (1.44), so we write down the probability that a fluctuation associated with the scale R_s is greater than δ_c:

$$p(\delta > \delta_c | R_s) = \frac{1}{\sqrt{2\pi\sigma(R)}} \int_{\delta_c}^{\infty} \exp\left(\frac{-\delta^2}{2\sigma^2(R_s)}\right) d\delta. \tag{1.45}$$

The Press-Schechter formalism now states that this probability is proportional to the probability that this point has ever been in a region with $\delta > \delta_c$. There is quite a subtle point in here, in that this assumes that any objects with $\delta > \delta_c$ are the ones which are just now undergoing gravitational collapse, i.e. it assumes $\delta = \delta_c$. If a point has $\delta > \delta_c$ then it would have $\delta = \delta_c'$, associated with a different mass and scale, M' and R_s', and would enter the mass function with that mass instead. This argument fails to account for underdense regions, and a factor of 2 is added to account for missed objects. We will accept this factor here, although an improvement on it can be found for example, in Peacock and Heavens (1990).

The mass function is then related to p as

$$\frac{Mn(M)}{\rho_0} = \left|\frac{dp}{dM}\right| \tag{1.46}$$

and we can write

$$n(M) = \frac{2\delta_c}{\sqrt{2\pi}} \frac{\rho_0}{M^2} \left|\frac{d\ln\sigma}{d\ln M}\right| \exp -\frac{1}{2}\frac{\delta_c^2}{\sigma^2}. \tag{1.47}$$

For a power-law mass fluctuations $\sigma(M) \propto M^{-\alpha}$ we get

$$n(M) \propto \left(\frac{M}{M_*}\right)^{\alpha-2} \exp\left[-\left(\frac{M}{M_*}\right)^{2\alpha}\right]. \tag{1.48}$$

Detailed numerical simulations give solutions which are different in detail, but the qualitative behaviour is correct. We find that the distribution of objects has a sharp cut-off at high masses, meaning large objects are more rare. Conversely, at the low mass end we have a shallower power-law slope. The Press-Schechter formalism has been extended and modified by subsequent work, which aim to find analytical routes that give more exact answers, and with them more insight (e.g. Lacey and Cole 1993, Sheth and Tormen 2002).

The analysis above is dominated by gravity, and appropriate for dark matter fluctuations which are governed by gravity interactions only. Nonetheless, we are now in a very good position - we have a recipe to identify and count dark matter virialised objects (dark matter halos) of any given mass in a density field evolved linearly. We now expect baryonic matter to form galaxies within the potential wells created by these objects (White and Rees 1978). However, it is gas dynamics and dissipative processes that shape the luminosity distribution of galaxies we *see* today.

If the mass to light ratio of galaxies was constant across luminosity or mass, the predicted luminosity function would be readily available from the mass function in (1.48), but this is far from being the case. White and Rees (1978) propose that galaxy formation is mainly regulated by how quickly gas is able to cool within a halo, which depends on its mass. Processes which regulate star formation within these gravitational wells are needed to explain the observed galaxy luminosity function, and the physics can get messy from now on. The observed galaxy luminosity function shows slopes for the high and low mass end which suggest that different star formation regulation mechanisms are at play at each extreme, measuring $\alpha - 2 \approx -1$ when fixing $2\alpha = 1$ (e.g. Bell et al. 2003). We will review some ways to tackle this problem in the next section.

Even though it is possible nowadays to detect the distribution of dark matter with weak gravitational lensing (Massey et al. 2007), the classical and easiest way to trace matter in the Universe is to map the luminous matter. The mission of constraining the matter power-spectrum would be easy if luminous matter was an unbiased proxy for matter, but this is not the case. This leads us to the concept of *galaxy bias*, which relates the luminous mass in the Universe, with the total amount matter which is present. The simplest case is a linear bias (Peacock 1999):

$$\Delta^2_{light} = b^2 \Delta^2_{matter}. \tag{1.49}$$

In reality this relation is likely to be more complicated, but the thing to keep in mind is that when we probe the galaxy population we are not directly probing the underlying density field and some assumptions are needed.

1.4.6 Galaxy formation models

Having reached this point we are now faced with the really difficult physics left to solve. We do not have, at this stage, a model for galaxy formation. This is largely due to the complexity of the system, rather than ignorance of the basic physical processes behind it. Our understanding of this highly complex process has been shaped by two different approaches.

One of them is sheer computational brute force. Ideally we would like to turn a set of potential wells and a primordial distribution of gas into a distribution of galaxies, by only inputting basic physics using *hydrodynamic* simulations (e.g. Pearce et al. 2001; Weinberg et al. 2004; Kereš et al. 2005). This is far from being within reach, and current simulations are limited in resolution and number of particles they work with. The upside is that computer power and numerical methods can only get better with time, and we expect this sort of simulations to give more accurate results as time goes on. In the meantime, processes which are beyond the resolution of the simulations have to be dealt with by analytical approximations.

A very different approach is to simulate only the gravitational interactions of dark matter, and use the resulting distribution as a starting point to a *semi-analytical* analysis (e.g. Kauffmann et al. 1999; Benson et al. 2003; De Lucia et al. 2006; Bower et al. 2006). Semi-analytical models rely on analytical approximations to complicated processes, such as star formation, gas cooling or feedback in order to predict a set of observables which can be matched to the real Universe. The advantage of semi-analytical modelling is that there is a very clear connection between the input and the output, which allows us to gain insight on which processes might be important in real galaxies.

Common between the two, and of particular relevance in this thesis, is their current inability to understand star formation. We also do not have a model for star formation in galaxies: given a galaxy of a given mass, luminosity or environment we are currently unable to predict what the star formation rate in that galaxy should be. Fundamental observables, such as the luminosity and stellar mass functions are shaped principally by star formation and merging. Given the theoretical difficulty in modelling these processes, observational constraints on how star formation depends on redshift, mass, luminosity, clustering, etc, are particularly crucial for the development of galaxy for-

mation models and consequently our own understanding of how they form and evolve. Many of these observables depend on an assumed initial mass function (IMF) - the predicted number of stars per unit mass formed from a single cloud of gas. An assumed IMF lies at the heart of any interpretation process relating to star formation within galaxies, as well as being an input for galaxy formation models.

This thesis aims to advance the current knowledge on this area by using the integrated spectra of galaxies. The spectrum of a galaxy holds vasts amounts of information about that galaxy's history and evolution. Finding a way to tap directly into this source of knowledge would not only provide us with crucial information about that galaxy's evolutionary path, but would also allow us to integrate this knowledge over a large number of galaxies and therefore derive cosmological information. These ideas are explored in Chapters 4 and 5.

1.5 The Cosmic Microwave Background as an observable

We have followed the evolution of primordial fluctuations in the inflation scalar fields to the stage where we can tentatively predict how the distribution and content of galaxies looks like today. In section 1.4.1 we mentioned how these primordial fluctuations $\delta\phi$ would also leave their signature in the radiation field. We now take a detour back in time, to look in more detail at how the observed Cosmic Microwave Background came to be, what it tells us about the Universe, and introduce some formalisms we will need in Chapter 2.

The CMB is an open window to the early Universe. It is a nearly-uniform and isotropic radiation field, which exhibits a measured perfect black-body spectrum at a temperature of 2.72K. This primordial radiation field is a prediction from a Big Bang universe - if in its early stages the Universe was at a high enough temperature to be fully ionized then processes such as Thompson scattering and Bremsstrahlung would thermalize the radiation field very efficiently. Assuming an adiabatic expansion of the Universe, one would then expect to observe a radiation field which would have retained the black-body spectrum, but at a much lower temperature.

As observers, we can measure three things about this radiation: its frequency spectrum $f(\nu)$, its temperature $T(\hat{\mathbf{n}})$ and its polarization states. Each of these observables

contains information about the creation and evolution of the field and are fully packed with cosmological information. Although the study of the polarization of the CMB radiation has been a recent and promising area of research (propelled by technology advancements which now allow this signal to be measured), this thesis concentrates on the temperature signal.

1.5.1 The CMB observables

The frequency spectrum of the CMB radiation was measured to high accuracy in the early nineties by FIRAS (as part of the COBE mission, which also gave us the first full-sky map of the CMB), and it was found to be that of a black-body at a temperature T=2.72K over a large range of frequencies. This profile indicates thermal equilibrium and it is to date the best example of a black-body known in the Universe. This alone can tell us something about the early Universe. If we assume an adiabatic expansion we expect $T \propto 1/R$. Relating the present day temperature to the temperature at a redshift z and using equation (1.7) we get

$$T_0 = \frac{T(z)}{1+z}. \tag{1.50}$$

This allows us to estimate the temperature of the radiation at the time the CMB was created. Our best estimate for the last scattering surface (LSS) redshift z_{LS} is approximately 1100, which gives us a temperature of around 3000K at the time of last scattering. And since $\nu_0 = \nu/(1+z)$, we expect a black-body spectrum to remain so in an adiabatic expansion (recall the flux of a black-body $B_\nu = \frac{2h\nu^3c^2}{e^{h\nu/KT}-1}$).

However, the vast majority of information lies not in the frequency spectrum of the CMB, but in its temperature field. Although the observed average temperature is amazingly uniform across the sky, a good signal-to-noise experiment will reveal small fluctuations around this average. These fluctuations are small (1 part in 10,000!), and in 2003 the satellite experiment WMAP provided the first high resolution, high signal-to-noise, full-sky map of these fluctuations. Since we are interested in the deviation from the average temperature, we generally define a dimensionless quantity $\Theta(\hat{\mathbf{n}}) = \frac{T(\hat{\mathbf{n}})-\langle T\rangle}{\langle T\rangle}$, where $\hat{\mathbf{n}}$ is a direction in the sky, $\hat{\mathbf{n}} \equiv (\theta, \phi)$.

We see these temperature fluctuations projected in a 2D spherical surface sky, and so it has become common in the literature to expand the temperature field using spherical

harmonics. The spherical harmonics form a complete orthonormal set on the unit sphere and are defined as

$$Y_{lm} = \sqrt{\frac{2\ell+1}{4\pi}\frac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(cos\theta)e^{im\phi} \tag{1.51}$$

where the indices $\ell = 0, ..., \infty$ and $-\ell \leq m \leq \ell$ and P_ℓ^m are the Legendre polynomials. ℓ is called the multipole and represents a given angular scale in the sky α, given approximately by $\alpha = 180/\ell$ (in degrees).

We can expand our temperature fluctuations field using these functions

$$\Theta(\hat{\mathbf{n}}) = \sum_{\ell=0}^{\ell=\infty} \sum_{m=-\ell}^{\ell} a_{lm} Y_{lm}(\hat{\mathbf{n}}) \tag{1.52}$$

where

$$a_{lm} = \int_{\theta=-\pi}^{\pi} \int_{\phi=0}^{2\pi} \Theta(\hat{\mathbf{n}}) Y_{lm}^*(\hat{\mathbf{n}}) d\Omega \tag{1.53}$$

and, analogously to what we do in Fourier space, we can define a *power spectrum* of these fluctuations, C_ℓ, as the variance of the harmonic coefficients

$$\langle a_{lm} a_{l'm'}^* \rangle = \delta_{\ell\ell'}\delta_{mm'} C_\ell \tag{1.54}$$

where the above average is taken over many ensembles and the delta functions arise from isotropy. We only have one Universe, so we are intrinsically limited on the number of independent m-modes we can measure - there are only $(2\ell+1)$ of these for each multipole. We can write the following expression for the power spectrum:

$$C_\ell = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} \langle |a_{lm}|^2 \rangle. \tag{1.55}$$

This leads to an unavoidable error in our estimation of any given C_ℓ of $\Delta C_\ell = \sqrt{2/(2\ell+1)}$: how well we can estimate an average value from a sample depends on how many points we have on the sample. This is normally called *cosmic variance*.

In real space, the power spectrum is related to the expectation value of the correlation of the temperature between two points in the sky:

$$\xi_{\Theta\Theta}(\theta) = \left\langle \Theta(\hat{\mathbf{n}})\Theta(\hat{\mathbf{n}}') \right\rangle = \frac{1}{4\pi} \sum_{\ell=0}^{\infty} (2\ell+1) C_\ell P_\ell \cos\theta, \;\; \hat{\mathbf{n}}.\hat{\mathbf{n}}' = \cos\theta. \tag{1.56}$$

1.5. THE COSMIC MICROWAVE BACKGROUND AS AN OBSERVABLE

Cosmological models normally predict what the variance of the a_{lm} coefficients is over an ensemble, so they predict the power spectrum. Each Universe is then only one realisation of a given model.

Under the Inflation paradigm, the temperature fluctuations are Gaussian, which means that the harmonic coefficients have Gaussian distributions with mean zero and variance given by C_ℓ. In this case, *all we need to characterise the statistics of our temperature fluctuations field is the power-spectrum* - higher-order correlation functions can be written in terms of the two-point function or the power spectrum. The Gaussian hypothesis is now being questioned by detections of non-Gaussianity and deviations from isotropy in the WMAP data. Part I of this thesis concentrates on testing the Gaussian hypothesis using the peaks of the temperature field.

The sum in equation (1.52) will generally start at $\ell = 2$ and go on to a given ℓ_{max} which is dictated by the resolution of the data. We exclude the first two terms for the following reasons: the monopole ($\ell = 0$) term is simply the average temperature over the whole sky ($Y_{00} = 1/2\sqrt{\pi}$ which makes $\Theta(\hat{\mathbf{n}})_{\ell=0} = 1/4\pi \int\int \Theta(\hat{\mathbf{n}}) d\phi dcos\theta \equiv \langle\Theta(\hat{\mathbf{n}})\rangle$, where the integrals are done over the entire surface), and so from our definition of $\Theta(\hat{\mathbf{n}})$ it should average to zero. The monopole temperature term would be a valuable source of cosmological information in its own right, but its value can never be determined accurately because of cosmic variance - essentially we have no way of telling if the average temperature we measure locally is different from the average temperature of the Universe. The dipole term ($\ell = 1, \alpha \approx 180^\circ$) is affected by our own motion across space - CMB photons that we are moving towards will appear blueshifted and those that we are moving away from will appear redshifted. This creates an anisotropy at this scale which dominates over the intrinsic cosmological dipole signal and therefore we normally subtract the monopole/dipole from a CMB map or discard the first two values of the power spectrum prior to any analysis.

Our best estimate at what the power spectrum of the observed CMB fluctuations looks like can be seen in Figure 1.3. It is usually plotted as $\ell(\ell+1)C_\ell/2\pi$. This is related to the contribution towards the variance of the temperature fluctuations in a patch of sky of size $\propto 1/\ell$: $\left\langle\Theta^2\right\rangle = \xi_{\Theta\Theta}(0) = \frac{1}{4\pi}\sum_\ell(2\ell+1)C_\ell$ (since $P_\ell(1) = 1$). The contribution over a range of values of ℓ is given approximately given by $\int_\ell^\infty 2\ell' C l' d\ell' = \int_\ell^\infty 2\ell'^2 C'_\ell \frac{d\ell'}{\ell'}$ (for $\ell \gg 1$) and so $2\ell^2 C_\ell$ is proportional to the contribution to the variance per unit $\ln \ell$.

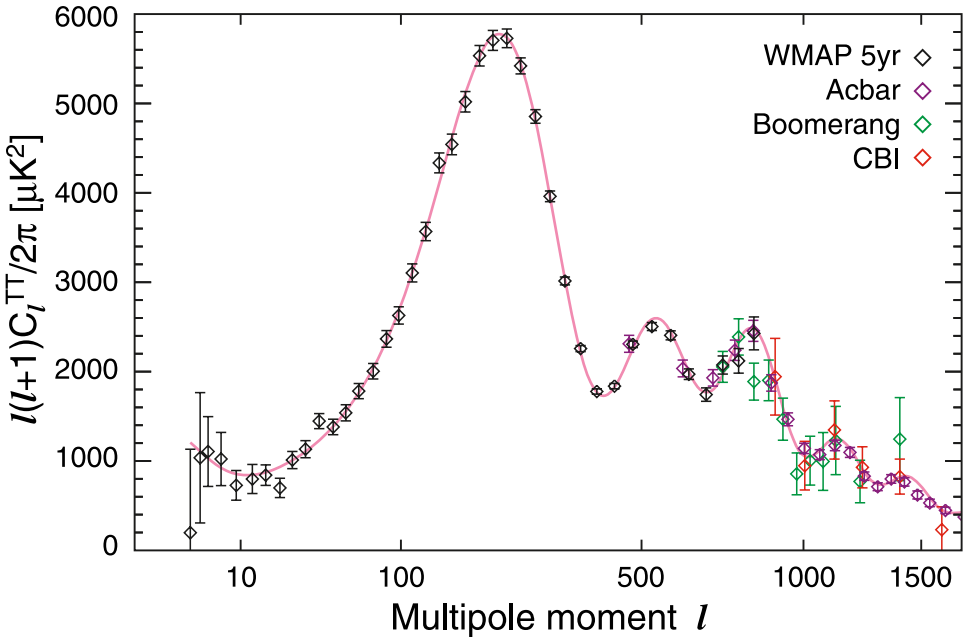

Figure 1.3: From Nolta et al. (2008). The CMB power spectrum as a function of angular scale, as derived from WMAP's five years of integrated data and three other small scale experiments: ACBAR (Reichardt et al. 2008), Boomerang (Jones et al. 2006) and CBI (Readhead et al. 2004). The red line is our best fit to the data.

This gives a flat plateau at large angular scales, and brings out a lot of the structure at smaller scales (see later).

1.5.2 Relating angular sizes with linear scales

It is useful to relate angular scales in the sky with linear sizes at the time of last scattering. We take the LSS as being a spherical surface at a redshift z_{LS} from us. We will take the comoving distance to this surface as being r_{LS}. We want to relate a small angle in the sky θ to the linear comoving distance x at last scattering, such that $\theta \approx x/r$ (for $\theta \ll 1$ and in flat space).

The comoving distance-redshift relation is given by equations (1.17) and (1.12). The integration can only be done numerically for most cases, but for the case of a matter-dominated, flat Universe then equation (1.12) simplifies to $H(z) = (1+z)^{3/2}$ and we get

$$r_{LS} = \frac{c}{R_0 H_0} \int_0^{z_{LS}} (1+z)^{-3/2} dz = \frac{2c}{H_0 R_0} \left(1 - (1+z_{LS})^{-1/2}\right). \tag{1.57}$$

For $z_{LS} \gg 1$ then r_{LS} is given simply by $\frac{2c}{H_0 R_0}$. Formally, by taking z_{LS} to infinity we are effectively calculating the present-day particle horizon length which is the maximum comoving distance light could have travelled since the Big Bang, d_H. So a small angle in the sky θ corresponds to a linear comoving distance x at last scattering given by (in radians)

$$\theta = x \frac{R_0 H_0}{2c} \tag{1.58}$$

One particular comoving distance at the time of last scattering which we might be interested in is the particle horizon length, which is given by

$$d_H(z = z_{LS}) = \int_{z_{LS}}^{\infty} \frac{dz}{H(z)} = \frac{2c}{H_0 R_0} (1+z_{LS})^{-1/2} \tag{1.59}$$

which, from (1.58) and for $z_{LS} \approx 1100$, means that

$$\theta_H^{LS} = (1+z_{LS})^{-1/2} \approx 1.7^\circ. \tag{1.60}$$

This tells us that scales larger than 1.7° in the sky were not in causal contact at the time of last scattering. However, the fact that we measure the same mean temperature across the entire sky suggests that all scales were once in causal contact. This was solved by the idea of Inflation, as introduced in section 1.4.1. In that section we saw how Inflation is a mechanism which provides us with primordial spatial inhomogeneities in the gravitational field, $\delta\Phi$, and with uniformity across the whole sky. We also saw how these inhomogeneities are the seeds of the large scale structure we see today. In the next section we will explore how they create the temperature fluctuations we observe in the CMB today.

1.5.3 Physical mechanisms: the origin of the anisotropies

CMB anisotropies can be classified into primary or secondary anisotropies, according to whether they were created at last scattering or during the photons' path along the line of sight. Photons can be affected by a range of things after last-scattering e.g. re-ionization, passing through hot clusters' gas, evolving potential wells, gravitational lensing, etc. While secondary anisotropies hold a good deal of information about the more recent Universe, they are not the subject of this thesis. Mostly, their effect on the temperature power spectrum lies at very small scales (very large ℓ) just beyond our current technical abilities. The exception is the Integrated Sachs-Wolf effect (related

to time-evolving potential wells) whose effect shows up at very large scales (very small ℓ), and causes a slight rise in the power spectrum.

Our interest in this thesis lies in the primary anisotropies. These, in turn, are created by two main mechanisms: gravitational and adiabatic

$$\Theta = \Theta_{grav} + \Theta_{ad}. \tag{1.61}$$

Perturbations in the gravitational potential $\delta\Phi$ left from Inflation can affect the radiation in two different ways. Firstly, through gravitational redshift, which in the weak field regime is given by

$$\frac{\delta\nu}{\nu} = \Theta \approx \frac{\delta\Phi}{c^2}. \tag{1.62}$$

Secondly, by causing a time dilation at the time of last scattering $\delta t/t = \delta\Phi/c^2$, which means we are looking at a younger universe when we look towards overdensities. In early times, $R \propto t^{2/3}$ and recalling that $T \propto 1/R$ we promptly get

$$\Theta \approx -\frac{2}{3}\frac{\delta\Phi}{c^2} \tag{1.63}$$

where we have again taken a weak-field approximation and assumed an adiabatic expansion. The added effect is simply

$$\Theta_{grav} \sim \frac{1}{3}\frac{\delta\Phi}{c^2} \tag{1.64}$$

which is commonly known as the Sachs-Wolf effect. These fluctuations happen at all scales, but dominate at large scales, where causal effects such as fluid dynamics (see next) do not come into account. For a spatial matter power-spectrum $P(k) \propto k^n$, the angular power-spectrum C_ℓ reduces to (for $n = 1$) $C_\ell \propto 1/\ell(\ell+1)$. This dependency gives rise to the flat part of the plot in Figure 1.3, which is usually called the Sachs-Wolf plateau.

We now turn to adiabatic perturbations. We have been slowly building up a picture the Universe at last scattering. Due to the high temperature, the Universe was fully ionized and consisted of a plasma mixture which, amongst others, contained photons and baryons. Thompson scattering meant that the photons were tightly coupled to the electrons which were in turn coupled to the baryons via Coulomb interactions. This coupling, together with radiation pressure acting as a restoring force, allows us to treat the primordial plasma as a perfect photon-baryon fluid to which normal fluid dynamics

equations apply.

As mentioned before, the Universe also displayed small local potential wells into which matter falls. These potential perturbations, $\delta\Phi$ can be related to matter density perturbations by Poisson's equation $\nabla^2\Phi = 4\pi G\rho_m$. For an adiabatic expansion, the matter density perturbations are related to the radiation density and temperature perturbations by

$$\frac{1}{3}\frac{\delta\rho_m}{\rho_m} = \frac{1}{4}\frac{\delta\rho_\gamma}{\rho_\gamma} = \Theta_{ad} \tag{1.65}$$

remembering that $\rho_m \propto R^{-3}$ and $\rho_\gamma \propto R^{-4}$.

We are interested in the dynamics of these temperature perturbations within this system. Let us take a simple model, in which we ignore gravity and the effect of the mass/intertia of the baryons (we are essentially taking a photon fluid), and see what happens to these temperature fluctuations under the influence of radiation pressure over time. The treatment we follow next is based on the excellent review by Hu and Dodelson (2002). We will be working in Fourier space: since the perturbations are small and evolve linearly we expect each k-mode to be independent.

The first thing to appreciate is that the number of photons is conserved. We can write down a continuity equation for photon number density, n_γ, as

$$\dot{n}_\gamma + \nabla.(n_\gamma \mathbf{v}_\gamma) = 0 \tag{1.66}$$

where the derivative is with respect to conformal time $d\eta \equiv cdt/R(t)$, which scales out the expansion, and $\mathbf{v}_\gamma$ is the photon fluid velocity. Taking into account the Universe's expansion, what is actually conserved is n_γ/R^3, and so $\dot{n}_\gamma + 3n_\gamma\frac{\dot{R}}{R} + \nabla.(n_\gamma\mathbf{v}_\gamma) = 0$ which reduces to

$$\left(\frac{\dot{\delta n_\gamma}}{n_\gamma}\right) = -\nabla.\mathbf{v}_\gamma \tag{1.67}$$

for linear perturbations $\delta n_\gamma = n_\gamma - \bar{n}_\gamma$.

We can relate this to temperature fluctuations by $n_\gamma \propto T^3$ to give $3\Theta = \delta n_\gamma/n_\gamma$. This reduces our continuity equation to $\dot{\Theta} = -(1/3)\nabla.\mathbf{v}_\gamma$ or, in Fourier space, to

$$\dot{\Theta} = -\frac{1}{3}i\mathbf{k}.\mathbf{v}_\gamma. \tag{1.68}$$

We now consider the momentum of the radiation. Momentum is given by $\mathbf{q} = (\rho_\gamma + p_\gamma)\mathbf{v}_\gamma$ where $(\rho_\gamma + p_\gamma)$ is the effective mass and, for radiation, $p_\gamma = (1/3)\rho_\gamma$. Ignoring gravitational effects and viscosities, the only force is given by the pressure gradient $\nabla p_\gamma = (1/3)\nabla\rho_\gamma$. We can then write $\dot{\mathbf{q}} = \mathbf{F}$ as

$$\frac{4}{3}\rho_\gamma \dot{\mathbf{v}}_\gamma = \frac{1}{3}\nabla\rho_\gamma \tag{1.69}$$

and so $\dot{\mathbf{v}}_\gamma = \nabla\Theta$ or, in Fourier space,

$$\dot{\mathbf{v}}_\gamma = i\mathbf{k}\Theta. \tag{1.70}$$

We now consider only the velocity component along the direction $\hat{\mathbf{k}}$, as this is the only one with a gravitational source and we write our final continuity and Euler equations as

$$\dot{\Theta} = -\frac{1}{3}kv_\gamma \text{ (Continuity)} \tag{1.71}$$

$$\dot{v}_\gamma = k\Theta \text{ (Euler)}. \tag{1.72}$$

These can quickly be combined to give

$$\ddot{\Theta} + \frac{1}{3}k^2\Theta = 0 \tag{1.73}$$

which is a simple harmonic oscillator equation. The 1/3 factor is generally the adiabatic sound speed which is defined as $c_s^2 \equiv \dot{p}_\gamma/\dot{\rho}_\gamma$ which in this case is equal to 1/3. The general solution for equation (1.73) is given by

$$\Theta(\eta) = \Theta(0)\cos(kc_s\eta) + \frac{\dot{\Theta}(0)}{kc_s}\sin(kc_s\eta). \tag{1.74}$$

By assuming negligible initial velocities and by defining a sound horizon as $s \equiv \int c_s d\eta$, we simplify our solution to $\Theta(\eta) = \Theta(0)\cos(ks)$.

Let us briefly summarise how far we have got. We are trying to analyse the dynamical behaviour of a photon-baryon fluid, and study how temperature fluctuations behave in this system. We took some very constraining assumptions (such as ignoring gravity and the baryons) and worked on a system whose only force was given by radiation pressure gradients. What we found is that this *pressure acts as a restoring force to initial perturbations and we are left with oscillations which propagate at the speed of sound*. This is an important result, which holds even when we take into account other effects to make our system a realistic one. This behaviour continues until we hit the

temperature of recombination, at which time matter and radiation de-couple and any temperature fluctuations are essentially frozen into the photons' temperature, which we measure (nearly unchanged!) today.

Let us emphasise that these oscillations are happening at all scales, and we are interested in those which, at the time of recombination, happen to be at one of their extrema. If this happens at a conformal time η_{rec} (corresponding to a sound horizon s_{rec}), then modes will be frozen with an amplitude given by

$$\Theta(\eta_{rec}) = \Theta(0)\cos(ks_{rec}) \tag{1.75}$$

and those caught at their extrema will have $k_n s_{rec} = n\pi$. We can therefore find a fundamental scale of oscillation by taking $n = 1$

$$k_F = k_1 = \frac{\pi}{s_{rec}}. \tag{1.76}$$

This is our largest oscillating mode, and of course, all of the corresponding overtones will be caught at their extrema too. These will correspond to higher values of k_n and are simply oscillations which have had time to go another complete half-oscillation: k_1 corresponds to the oscillation which has had time to compress fully once, $k_2 = 2k_1$ to the oscillation which has had time to compress and then decompress fully, and so on.

We see that the maximum scale at which these fluctuations will happen (related to $\frac{1}{k_F}$) is related to the sound horizon at the time of recombination, which was close to the particle horizon. This means that scales larger than this will not be affected by acoustic oscillations, and we would not expect otherwise given that acoustic oscillations can only happen in regions which are causally connected.

However, there is a caveat to this toy model. These oscillations also set up velocities in the fluid, which will in turn produce Doppler shifts in the frequencies of the photons. Velocity oscillations are precisely $\pi/2$ out of phase with acoustic oscillations which in this case cancels the temperature oscillations in the radiation completely and gives a flat Θ.

A full treatment should take into account gravity, mass and inertia of the baryons, the evolution of the photon/baryon ratio, viscosity, diffusion and so on. A full solution

looks more like (Hu 1995):

$$\langle \frac{\Delta T}{T}^2(\eta)\rangle = \left[\frac{1}{3}(1+3R)\Phi\cos(kc_s\eta) - R\Phi\right]^2 + \left[\frac{1}{3}(1+R)(1+R)^{-1/2}\Phi\sin(kc_s\eta)\right]^2 \tag{1.77}$$

where R is fluid momentum density and $R \equiv \frac{3\rho_b}{4\rho_\gamma} \approx \frac{450}{1+z}\frac{\Omega_b h^2}{0.015}$. The first term represents the acoustic oscillations which are equivalent to what we derived with our toy model and the second term represents the Doppler oscillations, which in this case are much smaller than the acoustic oscillations, but still $\pi/2$ out of phase (we notice the sine instead of the cosine).

However, we do not need a full treatment to understand what the effect of gravity and the introduction of baryons does to this system, at least qualitatively. Gravity's main effect is to introduce another force into the system, $\mathbf{F}_{grav} = -m\nabla\Phi$, which adds a term to our Euler equations. This in effect creates a potential well and changes the range of the temperature oscillations, effectively their amplitude. If we now introduce baryons to the system, then we are introducing mass into the system. As in a classical system consisting of a spring (restoring force) and a mass attached to the end of the spring (the photon-baryon fluid), increasing the mass will cause the spring to *fall* further, but it will not change the maximum rebound height. Recall that our peaks in the temperature fluctuations alternate between compression and expansion of the plasma, so introducing matter into the system changes only every other peak - those corresponding to compression of the fluid (matter is falling further into the potential well). Hence we expect even-numbered peaks to be suppressed in relation to odd-numbered peaks.

We also expect curvature to affect the observed angular temperature anisotropies, as it affects the path the CMB photons will have taken to get to us. It is essentially a problem of geometry: in a closed Universe, a given angle subtended in the sky will correspond to a smaller linear distance at last scattering than in a flat Universe, and so curvature shifts the peaks along the multipole axis. The detection of the first acoustic peak at an $\ell \approx 200$ provided a good constraint on the flatness of our Universe.

At smaller scales, however, we see that these oscillations are clearly damped. This comes from the fact that the last scattering surface has a finite thickness, and therefore recombination and last scattering do not happen at the same time. This damps out

small scale fluctuations (where the scale is related to the thickness of the scattering surface), as photons still have to random walk out of this shell before they are essentially free, smoothing out the fluctuations.

The WMAP satellite has launched us into the era of precision cosmology. Within the Gaussian hypothesis, all the information in the CMB is compressed in the power-spectrum and its analysis has revealed the most precise picture of the Universe to date - Table 1.1 summarises this picture. The extraction of cosmological information from the power-spectrum is a complex process in itself, and there are a number of degeneracies within the model that can only be lifted with the use of other datasets, such as type Ia supernovae or large scale structure surveys.

Parameter	Value
h	$0.719^{+0.026}_{-0.027}$
Ω_b	0.0441 ± 0.0030
Ω_c	0.214 ± 0.027
Ω_Λ	0.742 ± 0.030
Ω_m	0.258 ± 0.030
τ	0.087 ± 0.017
z_{dec}	1087.9 ± 1.2
z_{eq}	3176^{+151}_{-150}
n	$0.963^{+0.014}_{-0.015}$
σ_8	0.796 ± 0.036
t_0	13.69 ± 0.13 Gyr
z_{reion}	11.0 ± 1.4

Table 1.1: Summary of cosmological information derived from the analysis of the temperature power-spectrum, as estimated by WMAP5 (Komatsu et al. 2008). Parameters included on the table which have not been previously defined in this thesis are: τ as the optical depth at the time or recombination, t_0 as the age of Universe and z_{reion} as the redshift for reonization.

1.6 The integrated spectrum of a galaxy as an observable

In this section we will focus on the second observable used in this thesis - the optical spectra of a galaxy. Even though the stellar content of a galaxy is only the small tip of the iceberg, it remains a very important component of the Universe. Firstly because

we can see it, and secondly because it holds an imprint of that galaxy's star formation history, which combined with other galaxies' provides information of when, how and where luminous mass formed in the Universe.

Galaxies' integrated colours alone can provide insight about their evolution. The known bimodality of blue and red galaxies on a variety of observables seems to tell us that these two populations are intrinsically different. Whereas this is useful in its own right, there is a considerable amount of more information to extract from galactic light. Part II of this thesis concerns the problem of extracting information from a galaxy's integrated spectrum in a reliable way, and then using it to find out about the formation of structure in the Universe.

1.6.1 Stellar population models

First and foremost, this requires a means of physically interpreting galactic light. A galaxy's spectrum can be modelled as a superposition of stellar populations of different ages and metallicities, if we know the expected flux of each stellar population. This is given by stellar population models.

Single stellar population models (SSPs) have three main ingredients. First we need a description of the evolution of a star of given mass and metallicity in terms of observable parameters, such as effective temperature and luminosity (e.g. Alongi et al. 1993; Bressan et al. 1993; Fagotto et al. 1994a; Girardi et al. 1996; Marigo et al. 2008. This can be calculated (or at least approximated) analytically, to produce the so called isochrones: evolutionary lines for stars of constant metallicty in a colour-magnitude diagram. Secondly we need to assume an initial mass function (IMF), which gives the number of stars per unit stellar mass, formed from a single cloud of gas (e.g. Salpeter 1955; Chabrier 2003; Kroupa 2007). Different mass stars evolve with different time-scales, and we can use the IMF to populate different evolutionary stages of the colour-magnitude diagram with the correct proportion of stars of any given mass. Finally we need spectral libraries, which for a combination of parameters such as luminosity or colour index, assign a spectrum to a star. Spectral libraries can either be drawn from our local neighbourhood, by taking high quality spectra of nearby stars (Le Borgne et al. 2003), or they can be theoretically motivated (e.g. Coelho et al. 2007).

Stellar population models are limited in two main ways. Certain advanced stages of

stellar evolution, such as the supergiant phase, or the asymptotic giant-branch phase, are still poorly understood. This leads to uncertainties in the construction of the SSP models, which are in this case worsened by the fact that these are bright stars which contribute significantly to the overall luminous output. If using empirical spectral libraries, stellar population models are also limited by any bias of the stars in the solar neighbourhood. For example, the Milky Way is deficient in α-elements (O, Ne, Mg, Si, S, Ca, Ti), which are indicators of fast star formation. Nearby stars are biased towards low [α/Fe], which in turn bias the sample of high quality stellar spectra available for collection. In this case theoretical models might help, by explicitly calculating spectra for a variety of [α/Fe] models (Coelho et al. 2007).

1.6.2 Dust models

There is a further complication which arises from the fact that the light from each galaxy does not get to us without interference. Dust absorbs and re-emits light with a non-trivial wavelength dependence, both within each galaxy we want to observe and our own Milky Way.

For a uniform slab of dust, the emitted and observed flux are related by

$$F_\lambda^{obs} = F_\lambda^{em} e^{-\tau_\lambda} \tag{1.78}$$

where τ_λ is the optical depth of the obscuring material. This is clearly a simplification of the problem, and more sophisticated dust geometries can be found in Charlot and Fall (2000). These give a dependence on the optical depth which is more complex than the one in equation (1.78).

When talking about Galactic dust, it is more common to express the problem as a difference in magnitudes by writing

$$m_{\lambda,obs} - m_{\lambda,em} = -2.5 \log_{10} \left(\frac{F_\lambda^{obs}}{F_\lambda^{em}} \right) = 1.086 \tau_\lambda \equiv A_\lambda. \tag{1.79}$$

The difference in extinction in the B and V magnitude is called the colour excess defined as $A(B) - A(V) \equiv E(B-V)$. For a given extinction curve, k_λ, which holds the wavelength dependence of the problem, we generally write

$$A_\lambda = k_\lambda E(B-V) \tag{1.80}$$

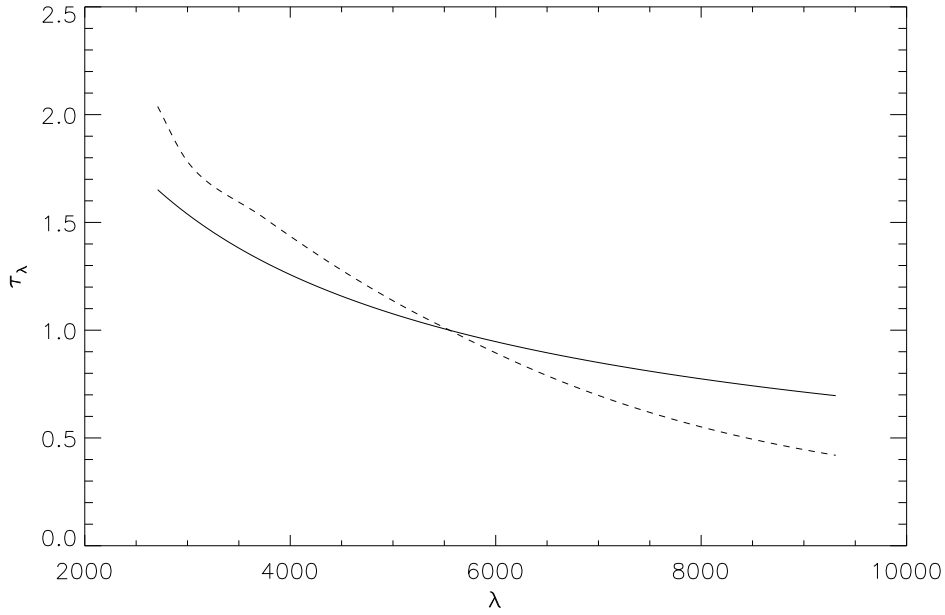

Figure 1.4: Two examples of dust extinction curves. The solid line shows a simple model that follows $\lambda^{-0.7}$ and is used throughout most of this thesis. The dashed line shows the the extinction curve estimated directly from the Large Magellanic Cloud by Gordon et al. (2003). Curves have been normalised to unity at $\lambda = 5550\mathring{A}$.

where the colour excess essentially provides the means for quantifying the amount of dust. k_λ can be theoretically or observationally motivated. Figure 1.4 shows the example of two absortion curves: one which simply goes as $\lambda^{-0.7}$ as mostly used in Charlot and Fall (2000) and in this thesis, and the extinction curve estimated directly from the Large Magellanic Cloud (LMC) by Gordon et al. (2003).

1.6.3 Extracting the information

Extracting information from galactic spectra is a much more complex problem than that of extracting information from, for example, the CMB's power-spectrum. Firstly we must be clear about the parameters we want to extract from the data. We are faced with a non-trivial decision, since any parametrization we might choose will undoubtedly be an over-simplification of the problem - a galaxy is almost infinitely more complex than the early Universe. However, the quality of the data will often impose a limit on how many parameters we can safely recover from the data and one must be careful not to ask for more than what the data allows. The risk is getting back a solution which is largely dominated by noise, rather than real physics.

From emission to absorption lines, continuum shape and spectral large scale features, a galaxy's spectrum is packed with information about the physics of that galaxy. Stellar population and dust models provide us with a theoretical framework for their interpretation, and there are various ways in which one can do this.

Certain isolated spectral features are known to be well correlated with physical parameters, such as mass, star formation rate, mean age, or metallicity of a galaxy (e.g. Kauffmann et al. 2003; Tremonti et al. 2004; Gallazzi et al. 2005; Barber et al. 2006). Emission lines are a sign of recent star formation: young, massive stars are the only ones with enough UV emission to ionize their surroundings. The recombination of the ionized gas creates signature emission lines, such as H_α and H_β, whose intensity (in the absence of dust) can tell us about the abundance of young stars in a galaxy. UV emission is, in itself, also a good probe for star formation for exactly the same reasons (e.g. Madau et al. 1996; Kennicutt 1998; Hopkins et al. 2000; Bundy et al. 2006; Erb et al. 2006; Abraham et al. 2007; Noeske et al. 2007; Verma et al. 2007).

Absorption features are directly related to the chemical abundances of a stellar population, as they are created when the black-body emission from the centre of the star passes through its cooler outer regions. Certain absorption features, such as the Lick indices, have been well measured and calibrated so as to provide a standard set of tools which aid in assigning a physical meaning to a given absorption line (e.g. Worthey 1994; Thomas et al. 2003).

This thesis focuses on using all of the available absorption features, as well as the shape of the continuum, in order to interpret a galaxy in terms of its star formation history. Emission lines are not included in the stellar population models (and are not present in every galaxy) and so we do not concentrate on these. We will show how, by using the integrated spectrum of a galaxy, we can find an appropriate parametrization which will allow us to recover the maximum amount of information from a galaxy without running into the risk of over-parametrizing.

1.7 Summary

We have seen how the origin and growth of the density fluctuations is connected to the distribution of galaxies we see today. The picture presented in this Introduction is normally referred to as the Standard Model of cosmology, or the Λ Cold Dark Matter

(Λ CDM) model. Even though this is perhaps the first time we stand with a cosmological model that describes the vast majority of our observations, there are thankfully many areas which leave our knowledge vastly unsatisfied. We have highlighted some throughout this introduction, but let us summarise the ones this thesis is particularly concerned with:

- What is the nature of Inflation? We are vastly ignorant about this crucial moment of the history of the Universe. There are very many proposed theoretical models, but without adequate observational constraints it is not possible to move theory forward. Inflation, if true, left its imprint on the statistics of the CMB, which we explore in Chapters 2 and 3.
- What is the nature of Dark Energy? Evidence that our Universe has recently entered an accelerated expansion state is clear in the luminosity-distance to type Ia supernovae. This exciting and direct probe of Dark Energy requires accurate knowledge of these explosions and their progenitor stars, for which we lack a theoretical understanding. We put observational constraints on the nature of these progenitors in Chapter 6.
- How do galaxies assemble their stellar mass? Our model predicts a hierarchical growth of galaxies, with smaller objects forming first and merging to form larger objects. We estimate robust stellar masses of a large number of galaxies and explore various issues relating to the formation and assembly of stellar mass in Chapter 5.

This thesis naturally splits into two parts because it focuses on two distinct fundamental observables of our Universe: the Cosmic Microwave Background, and the light from nearby galaxies. Both parts however, share the same goal: to further constrain or test our current model of cosmology and structure formation. We begin with non-Gaussianity studies of the CMB in Part I. Part II introduces VESPA, a novel algorithm which extracts information from a galactic spectra and which we then use to explore the local Universe.

Part I

Non-Gaussianity in the Cosmic Microwave Background

Chapter 2

Background and methodology

The hypothesis that the cosmic microwave background (CMB) is an isotropic Gaussian random field is a direct prediction from a large number of Inflation models. This Chapter describes how we have tested this hypothesis using the clustering properties of temperature peaks in the temperature field.

2.1 Background

According to the simplest scenarios, the initial conditions set by Inflation in the early-Universe produce Gaussian (or very nearly Gaussian) temperature fluctuations at the time of recombination. Testing the statistical property of Gaussianity in the observed CMB today therefore puts real constraints on the inflationary mechanism which laid down the primordial seeds of our Universe.

Testing the Gaussian hypothesis is not only of importance for its potential to tell us something about Inflation. Current analyses of the CMB, which aim to extract cosmological information from the observed temperature fluctuations, use the two-point correlation function (or angular power-spectrum in harmonic space) of the CMB as a compressed data-vector which statistically describes the underlying field completely. If the CMB is in fact Gaussian, then higher order statistics are null, and all the information is indeed accessible in the angular power spectrum. However, a significant detection of non-Gaussianity in the data could have important consequences not only

for our knowledge of the early-Universe, but also for our current cosmological model.

2.1.1 Searches for non-Gaussianity

It is therefore not surprising that a large number of studies have been dedicated to test the Gaussian hypothesis in the CMB. These searches started soon after the very first detection of the temperature fluctuations themselves with the Cosmic Background Explorer (COBE, Smoot et al. 1992), but it was not until the Wilkinson Microwave Anisotropy Probe experiment (WMAP, Bennett et al. 2003) that the data available were of enough angular resolution and signal-to-noise ratio to allow these studies to produce significant results.

Without a physically-motivated model for non-Gaussianity there are an infinite number of ways to modify a Gaussian random field so that it deviates from Gaussianity. It then becomes impossible to predict how sensitive a given estimator will be to a signal of unknown nature. This makes it advantageous to use a broad range of estimators, and CMB data has been analysed with an enormous array of statistics in real, harmonic and wavelet space.

The WMAP team have recently made their 3rd data release, corresponding to five years of integrated data. The first year data release very quickly yielded a large number of searches for non-Gaussianity (Colley and Gott 2003; Komatsu et al. 2003; Park 2004; Vielva et al. 2004; Coles et al. 2004; Copi et al. 2004; Cruz et al. 2005; Eriksen et al. 2004b,a, 2005; Land and Magueijo 2005a,b,c; McEwen et al. 2005; Mukherjee and Wang 2004; Gurzadyan et al. 2005; Liu and Zhang 2005; Tojeiro et al. 2006). Some of these studies reported anomalies, and some found the data consistent with Gaussianity. Most notably a north-south assymetry, an alignment of the low multipoles and a localised feature named the Cold Spot were found repeatedly by different teams and using different methods. Although some of these detections remained in subsequent data-releases (e.g. Cruz et al. 2007; Wiaux et al. 2008; McEwen et al. 2008), some have not (e.g. Dennis and Land 2008). It can be somewhat frustrating to conduct searches for non-Gaussianity that have no a priori physical mechanism behind them. A detection of a particular non-Gaussian feature more often than not offers little clue about its origin, even if we believe it to be cosmological. It is also instrinsically difficult to assess the significance of such detections, especially if we consider the infinite number of tests one

can conduct. Nevertheless, given a set of anomalies one can look for alternative models which might explain them. An example is a class of anisotropic cosmological models, the Bianchi models (Barrow et al. 1985), which have been investigated as a way to explain some of the CMB anomalies (e.g. Jaffe et al. 2005; McEwen et al. 2006). Even though there are problems reconciling these models with the concordance cosmological model, they demonstrate how one might learn from CMB anomalies.

2.1.2 f_{NL} models

Even within the inflationary paradigm, there is room for some level of non-Gaussianity if the evolution of the initial fluctuations from $\delta\phi$ to Θ is not completely linear. Specifically, there are three main possible sources: non-linearity in inflaton fluctuations $\delta\phi$, non-linearity in the $\delta\phi - \Phi$ relation and non-linearity in the $\Phi - \Theta$ relation.

It has become customary to parametrize all of these effects into one number as:

$$\Phi(\mathbf{x}) = \Phi_L(\mathbf{x}) + f_{NL}^{loc}(\Phi_L^2(\mathbf{x}) - \langle\Phi^2(\mathbf{x})\rangle) \tag{2.1}$$

where Φ_L is Gaussian and f_{NL}^{loc} parametrizes what in the literature is called the *local* non-Gaussianity. Single-field, slow-roll models of inflation predict f_{NL}^{loc} to be less than unity, whereas more complicated models predict much higher values of f_{NL}^{loc}. What does it mean to measure a positive value of f_{NL}^{loc}? Given that f_{NL}^{loc} parametrizes three potentially non-linear relations, it might not be immediately obvious which one causes the signal. However, models affect each stage of this evolution by different amounts, so a large value of f_{NL} would at the very least rule out a section of Inflationary models. Predicted experimental limits, based on the bispectrum, suggest that even an ideal experiment could only exclude the Gaussian hypothesis if $f_{NL}^{loc} > 3$, whereas WMAP and Planck require $f_{NL}^{loc} > 5$ and 20, respectively (Komatsu and Spergel 2001).

Finding the range of f_{NL}^{loc} values allowed by the data is therefore a way to *directly* differentiate between models. The most stringent constraint on f_{NL}^{loc} from the CMB comes from the analysis of the 5-year WMAP dataset by Komatsu et al. (2008) which finds $-9 < f_{NL}^{loc} < 111$ at a 95% confidence level. Finding a slighly broader constraint, but a more controversial central value, Yadav and Wandelt (2008) have estimated $27 < f_{NL}^{loc} < 147$ at the 95% level, therefore excluding the $f_{NL} = 0$ Gaussian hypothesis with high confidence.

The bispectrum (the Fourier transform of the three-point correlation function) is a widely-used estimator for f_{NL}^{loc} due to its high sensitivity to this type of non-Gaussianity (e.g. Komatsu and Spergel 2001). We can write

$$\langle \Phi(\mathbf{k_1})(\mathbf{k_2})(\mathbf{k_3}) \rangle = (2\pi)^3 \delta^3(\mathbf{k_1} + \mathbf{k_2} + \mathbf{k_3}) F(k_1, k_2, k_3) \tag{2.2}$$

where F is the three-point function. For the type of non-Gaussianity mentioned above, F is large for configurations that have $k_1 \ll k_2, k_3$. Recently, another type of non-Gaussianity has been introduced in which F is large for equilateral configurations: i.e. $k_1 \sim k_2 \sim k_3$. With this type of non-Gaussianity we associate f_{NL}^{equil}. WMAP5 constraints on this type of non-Gaussianity: $151 < f_{NL}^{equil} < 253$ at the 95% level (Komatsu et al. 2008). The advantage is that, between the two types on f_{NL} models, they are sensitive to most inflationary models proposed in the literature.

Non-Gaussianity in Φ will also propagate into the distribution of galaxies. There is a further complication here, in that the distribution of galaxies is affected by non-linearities both from gravitational evolution and galaxy bias (Sefusatti and Komatsu 2007). So to probe primordial non-Gaussianity using large-scale structure, we first must understand these. The returns, however, are high. The scales probed by large scale structure are smaller than those present in the CMB, and the two measurements are complementary. Sefusatti and Komatsu (2007) have shown that an all-sky survey up to redshift 5 would detect $f_{NL}^{loc} \sim 0.2$ and $f_{NL}^{equil} \sim 2$, which is a vast improvement on CMB hopes. Slosar et al. (2008), using a formalism introduced by Dalal et al. (2008), combine the WMAP5 dataset and large scale structure to estimate $-29 < f_{NL}^{loc} < 69$ at the 95% level.

2.1.3 Interpretation

The difficulty in interpreting a detection of non-Gaussianity goes beyond a sound statistical analysis for two reasons. Firstly, secondary anisotropies such as the Sunyaev-Zel'dovich (SZ) effect (Sunyaev and Zeldovich 1972), the Ostriker-Vishniac effect (Ostriker and Vishniac 1986; Castro 2003) or the Rees-Sciama effect (Rees et al. 1968), created along the line-of-sight, act on the CMB photons in a way which produces deviations from non-Gaussianity. Packed with cosmological information in their own right,

these secondary anisotropies manifest themselves at too small scales to be measurable in the WMAP data. The exception is the integrated Sachs-Wolfe (ISW) effect, which is significant at scales $\ell < 30$. Francis & Peacock (2008; in prep) reconstruct the local density field for $z < 0.3$, which they then use to estimate the ISW contribution to the observed CMB. Removing this contribution therefore allows a better estimation of the cosmological signal at large scales.

Secondly, Galactic and extragalactic emission present a serious problem. Highly non-Gaussian in nature, foreground emission needs to be carefully separated from the CMB signal prior to analysis. The WMAP team has produced foreground-cleaned maps for three of the observed frequencies (see section 2.2.1). Nevertheless, it is still unclear how much residual foreground power can be found in the temperature maps, and in what way this power might affect a given estimator. The potential existence of un-removed foregrounds and systematic instrumental effects not accounted for in the data-reduction pipelines remain the usual suspects for any detection of non-Gaussianity, although it is also normally not possible to completely exclude a cosmological origin.

This section focuses on using the two-point correlation function of temperature maxima and minima (hot and cold spots) as a probe of non-Gaussianity. The two-point correlation function of peaks of a Gaussian random field depends only on its power-spectrum. In brief, we use Gaussian simulations of the CMB with the same observed power spectrum as the observed CMB and compare the resulting peak-peak auto- and cross-correlation functions to that estimated directly from the data. Using non-Gaussian simulations with $f_{NL}^{loc} \neq 0$, we also explore how sensitive peak statistics are to this particular type of non-Gaussianity. The rest of this chapter describes the data (in Section 2.2) and methodology used to generate the Gaussian maps (in Section 2.3.2) and estimate the peak statistics (in Section 2.3.5). This method was applied to the first and the third WMAP data releases, and the results are presented in the next Chapter.

2.2 Data

The WMAP satellite probed the CMB at five different frequencies with two radiometers, producing ten differencing assemblies (DAs): four on the W-band (94GHz), two on the V-band (61GHz), two on the Q-band (41GHz), one on the Ka-band (33GHz)

and one on the K-band (23GHz). Each of these assemblies, after calibration and removal of the monopole and dipole, is available for download[1]. The satellite has been in operation for over five years, producing three data releases corresponding to one, three and five years of integrated data.

All the maps are provided in the Hierarchical Equal Area isoLatitude Pixelisation (HEALPix) scheme[2], which has proved to have several advantages over other methods for pixelising the surface of a sphere, in particular the fact that the pixel area is kept constant throughout the surface of the sphere. However, the pixel shapes can vary largely between the equatorial and polar regions and distance between pixel centres is not kept constant. The HEALPix scheme divides the sphere surface into 12 faces of 4 sides each, giving a minimum resolution of 12 pixels. Each side is divided in N_{side} pixels, giving a total number of pixels in a map of $12N_{side}^2$. The WMAP maps were provided at a resolution of $N_{side} = 512$ giving a total of 3,145,728 pixels separated on average by $\theta_{pix} = 0.115$ degrees $= 6.87$ arc minutes.

Each DA map pixel p contains the temperature field (in mK) and a field containing the number of observations, $N_{obs}(p)$, which allows the noise per pixel to be estimated using

$$\sigma(p) = \frac{\sigma_0}{\sqrt{N_{obs(p)}}} \tag{2.3}$$

where σ_0 is the noise dispersion per map and which has been published for each of the different assemblies (Bennett et al. 2003). Also available is a foreground-cleaned map of each of the DAs (see 2.2.1), from which a Galactic foreground template has been removed, consisting of synchrotron, free-free and dust emission; and the beam transfer functions of each receptor $b(\ell)$ from which the corresponding window function $W(\ell)$ can be calculated ($W(\ell) = b^2(\ell)$).

2.2.1 Foreground emission

As mentioned in section 2.1, galactic and extra-galactic radio emissions need to be identified and subtracted from the observed data. Radio emission from galaxies arises mainly from three mechanisms: non-thermal synchotron emission from relativistic electrons which spiral along the lines of large-scale magnetic fields; free-free (bremsstrahlung)

[1] http://lambda.gsfc.nasa.gov/product/map/m_products.cfm
[2] http://www.eso.org/science/healpix/

emission from the decceleration of electrons by charged particles; and thermal dust emission in the radio band. The first two dominate at lower frequencies and decline with frequency, whereas the latter has the opposite behaviour. Figure 2.1 sketches the relative contributions of each of these mechanisms against the CMB power.

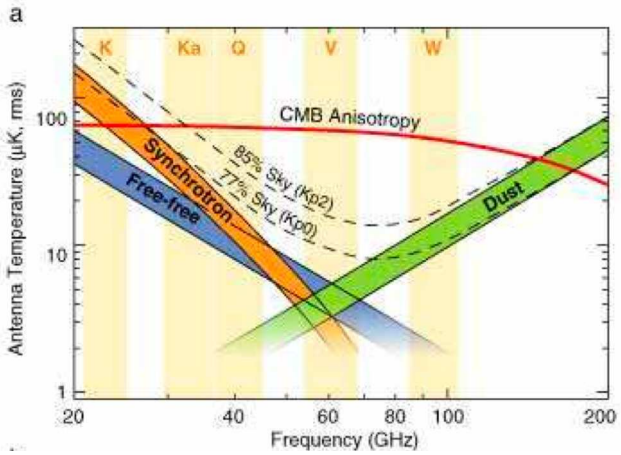

Figure 2.1: From Bennett et al. (2003): contribution of diffuse radio emission mechanisms from within the Milky Way and their frequency behavious in comparison with the CMB signal.

The problem of separating the foreground from the CMB signal can be tackled with two different approaches. One relies simply on CMB data, and combines CMB maps at different frequencies with weights chosen as to minimise the temperature variance on a pixel-by-pixel basis. Examples are the Internal Linear Combination (ILC) map from Bennett et al. (2003) and the Tegmark Cleaned Map (TCM) from Tegmark et al. (2003).

The other approach is to explicitly construct models for the frequency dependence of the radio emission of each of the foreground components, and use a fitting technique to construct a model of the galactic emission. Once constructed, it can then be removed from the data. This approach has the additional advantage of providing insight into the physical mechanisms involved, although it often relies on outside data to construct template maps as starting points for the fitting processes. Galactic templates have been constructed for the WMAP data for all of its data releases (Bennett et al. 2003; Hinshaw et al. 2007; Gold et al. 2008), and subtracted from each DA to provide foreground-cleaned maps at for the 8 radiometers in the Q-, V- and W-band. This approach is more flexible from a user point of view, as it allows us to construct foreground-cleaned maps at any frequency, and to linearly combine different frequencies

in order to boost the CMB signal (see section 2.2.4).

On small scales, the problem lies not with diffuse Galactic emission but with extra-galactic point sources. These can be estimated by looking for bright temperature peaks which approximate the beam profile, and cross-matching with known radio sources (see Wright et al. 2008 for full details). Some sources will be missed due to being faint, but their contribution to the power-spectrum can be estimated (e.g. Komatsu and Spergel 2001; Komatsu et al. 2003, using the bispectrum).

Even after a foreground model has been removed some pixels still contain a large contribution from Galactic emission. Sky masks have been produced by the WMAP team which flag such pixels allowing the user to exclude them from a CMB analysis. Depending on the level of residual foreground deemed acceptable, several masks have been produced and made publicly available. The masks kp2 and KQ85 have been suggested as appropriate by the WMAP team for cosmology analysis of the CMB, for the 1st and 3rd data releases respectively.

2.2.2 Instrumental systematic effects

There are a vast number of instrumental systematic effects which need to be taken into account before interpreting and simulating CMB data. A comprehensive description and treatment can be found in Hinshaw et al. (2003a). Here we will briefly discuss two of the aspects which impact directly on the task of simulating a CMB sky.

The WMAP is a differencing instrument, which measures the difference in temperature between two parts of the sky. Whereas this is easy to do for two points observed at nearby times, time-drifts of the instrumental and other background sources means that it can be hard to do when the points are observed some time apart. In practice, it leads to an additional noise term which is commonly called $1/f$ noise (due to the typical frequency dependence observed in this type of noise). $1/f$ noise becomes important at large scales and it presents a deviation from the white-noise spectrum which dominates the experiment at smaller scales. Figure 2.2 shows the power-spectrum of simulated noise with in-flight properties compared to the power-spectrum of the CMB for each of the eight DAs. Fortunately, the signal dominates at the scale at which $1/f$ noise could be a problem, and we do not have to worry about it when simulating Gaussian maps,

with few exceptions (see Section 2.3.2).

The CMB observed with each of the WMAP radiometers is a convolution of the true signal with the window function of each of the receptors. It therefore becomes essential to have an accurate characterisation of the beam transfer functions. This was done using observations of the planet Jupiter (Page et al. 2003; Jarosik et al. 2007; Hill et al. 2008) to construct beam profiles, which in turn were used to estimate the beam transfer functions. The third- and fifth-year data releases introduced improvements to the modelling of the beam profiles, as more observations became available.

2.2.3 Estimating the temperature power-spectrum

In the ideal case where the CMB can be observed free of noise and foregrounds across the sky, and in which the window function is known exactly, the temperature power-spectrum C_ℓ as defined in equation (1.55) can be easily calculated by estimating the harmonic coefficients $a_{\ell m}$, using equation (1.53) directly from the data. The only limitation would be cosmic variance. However, in the real world the function $T(\hat{\mathbf{n}})$ is not known exactly. At low-ℓ we have the inevitable presence of foregrounds as discussed in 2.2.1, which often lead to cuts across the sky. At high-ℓ the noise dominates over the signal.

The exact methodology employed by the WMAP team to estimate the temperature power spectrum evolved over the three data releases. In the first year analysis, they used a quadratic estimator that computes a pseudo-power spectrum, $\tilde{C}_\ell$, from a masked map using essentially equation (1.53). Even though the pseudo-power spectrum is clearly different from the true one, their ensemble averages can be associated via a coupling matrix which itself depends on the form of the window function (see Hinshaw et al. (2003b), Appendix A for full details). An advantage of this method is that it allows $\tilde{C}_\ell$ to be estimated from the cross-correlation of two different DAs, replacing equation (1.53) by

$$\tilde{C}_\ell = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} \tilde{a}^i_{lm} \tilde{a}^{j*}_{lm}. \tag{2.4}$$

If the noise between the two DAs is uncorrelated, then the estimation of the pseudo power spectrum is independent of the noise in any given channel.

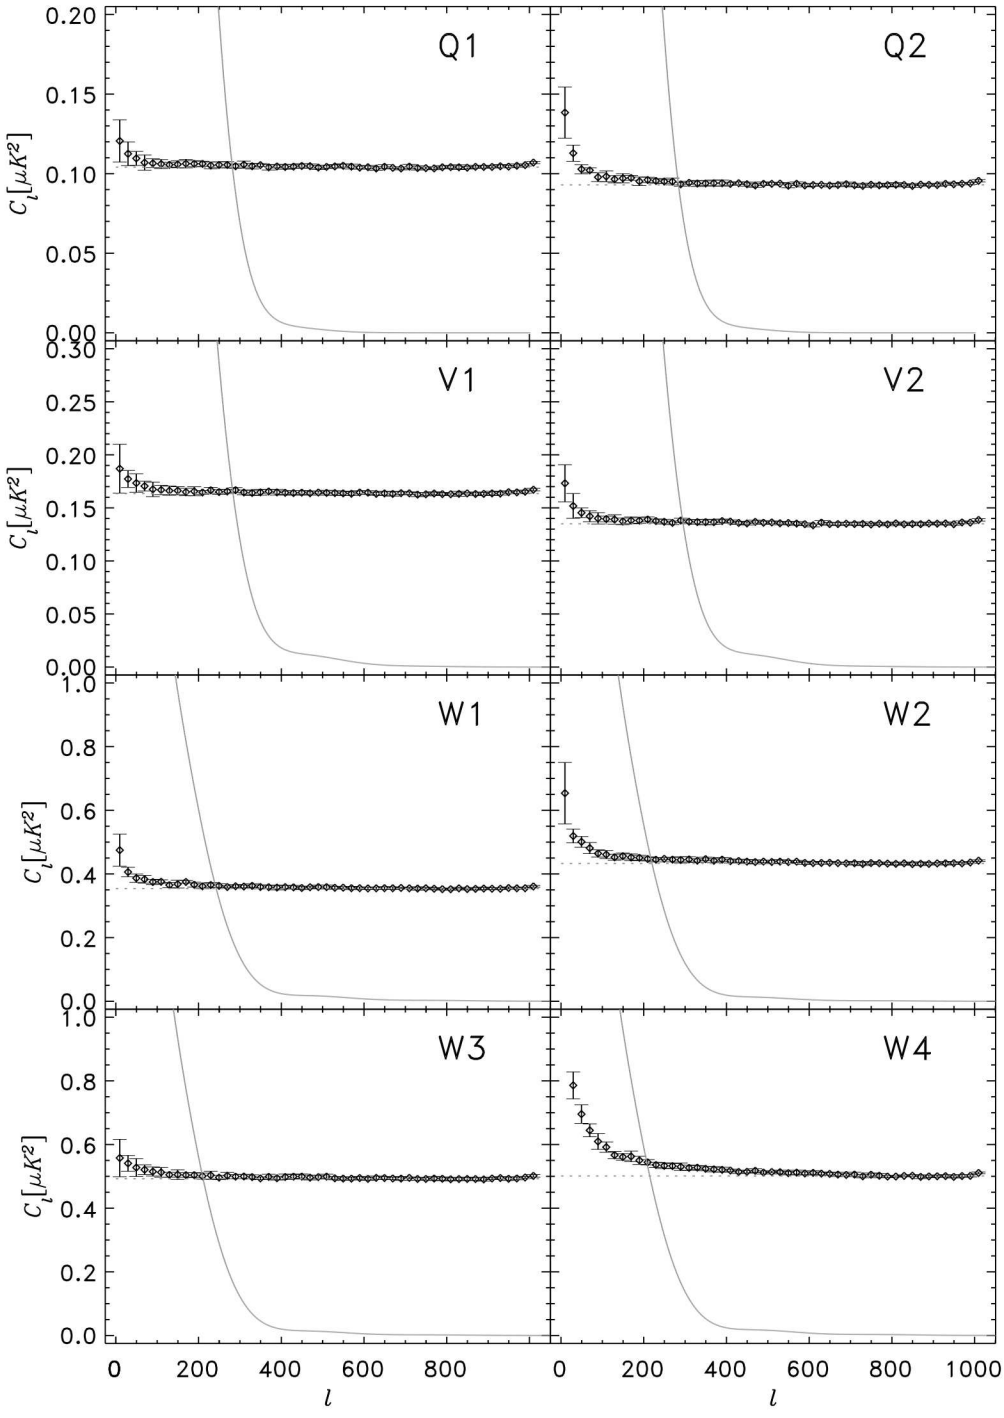

Figure 2.2: From Hinshaw et al. (2003b): the power-spectrum of noise as a function of ℓ, simulated with a full end-to-end pipeline and including $1/f$ noise is shown in the black line. The grey line shows the power-spectrum of the temperature fluctuations. At large scales, where the $1/f$ noise becomes important, the signal dominates over the noise and at small scales the noise is white and uncorrelated.

In the subsequent two data releases this method was only used for small scales, and the WMAP team adopted other methods to estimate C_ℓ for $\ell < 30$. One of them relies simply on using a full-sky ILC map to estimate the power spectrum directly. The other option is to use a maximum likelihood estimate, which estimates the best-fit power spectrum by maximising the probability of C_ℓ given the observed data (details in Hinshaw et al. (2007)). This gives the optimal estimate of C_ℓ, but it is computationally expensive.

The WMAP team has settled for a hybrid method for power spectrum estimation, in which the large scale power is estimated using a maximum likelihood estimate and the power at small scales is computed using a quadratic estimator.

2.2.4 CMB maps

To increase the signal-to-noise ratio of the data, we use linear combinations of the foreground-cleaned assemblies in the Q, V and W bands. We combine the maps in real space, at the original resolution, which keeps the noise uncorrelated at small scales. To combine two or more assemblies (e.g. X and Y) we calculate, for each pixel p:

$$T_{XY}(p) = \frac{\sum_{j=\{X,Y\}} T_j(p) w_j(p)}{\sum_{j=\{X,Y\}} w_j(p)} \tag{2.5}$$

with the weights being given by

$$w_j(p) = \frac{1}{\sigma_j^2(p)} \tag{2.6}$$

The index j corresponds to the different DAs: $j = 1, 2$ corresponds to the V band, $j = 3, 4$ to the Q band and $j = 5$ to 8 to the W band. This allows us to construct co-added or single frequency maps.

In spite of the foreground cleaning process no map is suitable for a full-sky analysis of the CMB and we mask all maps prior to analysis (see section 2.2.1).

2.3 Methodology

Our approach is to simulate an ensemble of Gaussian and non-Gaussian maps with the observed temperature power spectrum of the CMB and instrumental properties of the

WMAP satellite. This allows us to estimate the two-point peak-peak auto- and cross-correlation functions and their variance from the simulations, and compare it with that estimated directly from the data.

The methodology used can be separated into three main parts: the construction of simulated CMB maps (sections 2.3.2, 2.3.4 and 2.3.3), the estimation of the peak-peak correlation function from the maps (section 2.3.5), and the statistics used to analyse the results (section 2.3.6). We start by summarising some HEALPix routines, which were extensively used whilst doing this work.

2.3.1 HEALPix

In conjunction with defining a pixelisation scheme for the surface of the sphere, HEALPix provides a software suite to act on the maps both in real and harmonic space and perform a series of useful operations:

- `hotspot`: returns the position and value of all local extrema in a map,
- `anafast`: returns the harmonic coefficients a_{lm} of a map,
- `synfast`: generates a map in real space, given either the collapsed a_{lm} coefficients or a power-spectrum $C(\ell)$,
- `smoothing`: convolves a map with a gaussian beam,
- `ud_grade`: changes the resolution of a map.

2.3.2 The Gaussian maps

To construct Gaussian simulations of the CMB, we follow the method suggested by Komatsu et al. (2003) and proceed in the following way:

1. We generate one sky realisation from the best fit ΛCDM model power spectrum, published for each data-release, using synfast.
2. We copy this map n times, one for each assembly, and convolve each of the copies with the appropriate window function, using the harmonic coefficients extracted with anafast.
3. We add uncorrelated noise to each of the maps in real space, according to equation (2.3) (a more accurate noise model is used for difference maps).

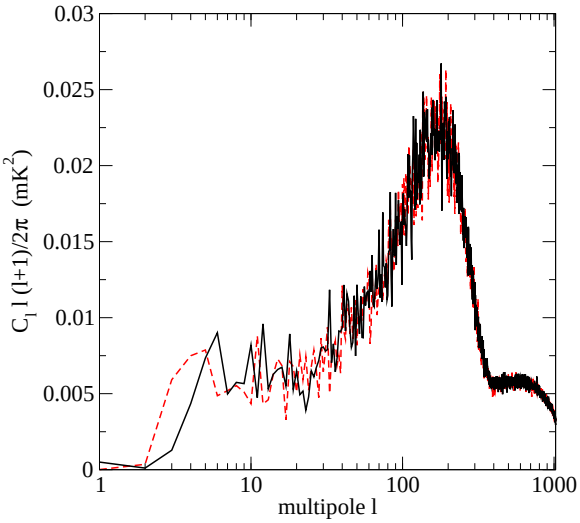

Figure 2.3: The power spectrum of our working maps. WMAP data in the black line and one of our Gaussian maps in the dashed red line.

4. We combine the n resulting maps using equations (2.5) and (2.6).

We repeat this procedure to create many Gaussian simulations of the CMB, each being a random Gaussian realisation of the same initial power spectrum. We used different numbers of Gaussian maps in different types of analysis, and we quote each number within the appropriate section. The maps are time-consuming to produce but in each case we check convergence of χ^2 (see Figures 3.2, 3.7 and 3.9 for examples). A comparison of the power spectrum of the real and a simulated map can be seen in Figure 2.3.

Although at small angular scales the noise properties are white, fully understood and easily modelled, at large angular scales individual Q, V and W assemblies present noise characteristics which are non-white. Fortunately these are entirely dominated by the signal and one does not need to worry about them (section 2.2.2). The WMAP team have produced a set of 110 noise maps which include white noise (dominating at small scales), $1/f$ noise (dominating at large scales) and inter-channel correlations for each of the radiometers. Ideally one would like to incorporate all known effects into the analysis. However, being limited by the relatively small number of full noise simulations and due to the high signal-to-noise ratio at the scales where the noise properties deviate from white, we choose to include only white noise in our Gaussian co-added and single frequency maps.

2.3.3 High-pass filtered maps

In order to limit the effect of cosmic variance on the peak statistics (see section 2.3.5) we have also analysed high-pass filtered maps, which remove the large-scale signal. This operation was performed on real, Gaussian and non-Gaussian maps.

We constructed several window functions given by $\mathcal{W}_{\ell_{cut}}(\ell) = 0$ for $\ell \leq \ell_{cut}$, $\mathcal{W}_{\ell_{cut}}(\ell) = 1$ otherwise. We mask the real maps before filtering. This is necessary because of the presence of foregrounds - the strong ringing effect in pixel space which results from such a sharp cut-off in harmonic space causes unwanted foreground signal to leak from the masked region. We follow the algorithm described below:

- We mask the WMAP data,
- We convolve the map with $\mathcal{W}_{\ell_{cut}}$, using the harmonic coefficients extracted by `anafast` from the masked map,
- We generate the map in pixel space using `synfast`,
- We re-mask the map and remove any residual monopole/dipole from the unmasked regions.

Since there is no foreground contamination in the Gaussian maps, there is no need to apply the initial mask. For testing purposes, we applied both methods to a number of Gaussian maps and found them to produce the same final results.

2.3.4 Non-Gaussian maps

In our analysis of the fifth-year data, we use non-Gaussian maps with $f_{NL}^{loc} \neq 0$ to estimate the sensitivity of peaks statistics to f_{NL}^{loc}. We will not consider f_{NL}^{equi} for the rest of this thesis, so we will drop the *loc* subscript. We use high-resolution ($\ell_{max} = 2901$, $N_{side} = 1024$) f_{NL} maps to simulate the 70GHz and the 100GHz bands of the upcoming Planck satellite.

An algorithm to generate non-Gaussian maps following equation (2.1) is given by Liguori et al. (2003) This generates Gaussian and non-Gaussian harmonic coefficients (a_{lm}^{G} and a_{lm}^{NG}, respectively) which can be used to construct the coefficients of a given map by

$$a_{lm} = a_{lm}^{G} + f_{NL} a_{lm}^{NG}. \tag{2.7}$$

We assume perfectly Gaussian beams and anisotropic noise (the number of observations per pixel was given by the Planck Sky Model), and use the predicted instrumental properties. We mask the maps using the WMAP kp0 mask and smooth with a FWHM=10 arcseconds beam.

The power-spectrum of these maps is matched to that of the data we wish to analyse. For each set of a_{lm}^G and a_{lm}^{NG} we generate one sky realisation, and produce a non-Gaussian map following the steps 2-4 described in section 2.3.2.

2.3.5 Estimating $\xi(\theta)$

There are several estimators suggested in the literature to estimate $\xi(\theta)$ directly from the data. They all work by comparing the sample of points to an uniform, random catalogue with the same spatial distribution as the real data. We used the Hamilton (1993) estimator, which promises fast convergence:

$$\xi(\theta) = \frac{RR(\theta).DD(\theta)}{DR(\theta)^2} - 1 \tag{2.8}$$

where $RR(\theta)$ and $DD(\theta)$ are the number of random and data pairs respectively at a distance θ from each other and $DR(\theta)$ is the number of cross-pairs separated by a distance θ (all weighted by the number of total random, data and cross pairs in the catalogue). Indeed, we found it to converge faster than the standard estimator, $\xi(\theta) = \frac{DD(\theta)}{DR(\theta)} - 1$. We use large random catalogues with the same sky cut as the appropriate WMAP map, and ensure that the estimator has converged to a stable value. A hot spot (cold spot) is defined for the purposes of this analysis as the centre of any pixel whose temperature is higher (lower) than the temperature of all pixels with which it shares a boundary.

In our analysis of the third-year data, we also explore the use of the cross-correlation function of peaks between different maps of different frequencies as a probe of non-Gaussianity. If the noise is uncorrelated from one detector to the next we should expect a higher sensitivity to real temperature peaks by cross-correlationg the measurement from two frequencies. Again we use the Hamilton estimator, modified do as to take into account two independent sources of peaks (Mann et al. 1996):

$$\xi(\theta) = \frac{AB(\theta)RR(\theta)}{AR(\theta)BR(\theta)} - 1 \tag{2.9}$$

where the pair counting is defined as above, with the difference that instead of one we have two data catalogues, corresponding to the letters A and B.

Both estimators are calculated using simple pair-counting methods. The code was tested by comparison with the analytic prediction from Heavens and Gupta (2001) in the ideal case of an unmasked sky without the addition of noise.

We consider all temperature hot spots above $\nu\sigma$ and cold spots below $-\nu\sigma$. There is no a priori reason to choose any particular temperature threshold. As we increase the threshold the number of peaks decreases and we are limited by cosmic variance - the amplitude of the large scale multipoles can easily change the number of hot and cold spots above a given threshold. As we decrease the threshold the number of peaks increases and the calculation becomes computationally prohibitive. We want to choose the value of ν which allows us to analyse the most number of peaks within a reasonable time scale. In our first year analysis this value was $\nu = 1.5$, which as we will see leaves us limited by cosmic variance. Given improvements in hardware and computer code, we were able decreased this value down to $\nu = -1$ in our fifth-year analysis.

However, increasing the number of peaks might not be the most appropriate thing to do when searching for f_{NL}. At large scales for example, we know the temperature fluctuations are dominated by the Sachs-Wolfe effect (equation 1.64) and $\Theta \propto \delta\Phi$. Adding a non-linear component to the potential then only changes the amplitude of the temperature fluctuations relatively to the Gaussian case, but not the peaks' positions. This in fact suggests that on large scales, we might increase the sensitivity to f_{NL} by selecting threshold which selects only some of peaks, as for each value of f_{NL} we expect this threshold to select a different set of peaks. The optimum choice of threshold is a balance between noise, computational time and sensitivity, and for each case it can only be found empirically right now.

2.3.6 Statistics

Testing the Gaussian hypothesis

We use the χ^2 statistic to interpret our results. For each map we calculate

$$\chi^2 = \sum_{i,j} (\xi_i - \bar{\xi_i}^G) C_{ij}^{-1} (\xi_j - \bar{\xi_j}^G) \tag{2.10}$$

where the covariance matrix C_{ij} and the mean values $\bar{\xi}^G$ are estimated from the Gaussian maps. i and j identify bins at a given angular separation. The results will be presented in terms of the reduced χ^2, obtained by dividing χ^2 by the number of de-

grees of freedom. In our case, this is simply the number of points used to evaluate equation (2.10). We specify this number at each relevant section.

We compare the value of χ^2 obtained from the observed CMB map with the one-point distribution function of the values of χ^2 obtained from an ensemble Gaussian maps in order to interpret the significance of a detection. We should note, however, that any significance that is estimated in this way is likely to be over-estimated. A full treatment should take into account the total number of independent non-Gaussianity tests performed in any given map.

We also note that we do not assume Gaussianity when assigning confidence levels to any χ^2 value. We use χ^2 as a statistic whose probability distribution is empirically estimated from Monte Carlo simulations.

Constraining f_{NL}

When trying to constrain f_{NL} we find the minimum of

$$\chi^2(f_{NL}) = \sum_{i,j} (\xi_i - \bar{\xi}_i(f_{NL}))C_{ij}^{-1}(\xi_j - \bar{\xi}_j(f_{NL})) \tag{2.11}$$

with respect f_{NL}. The mean values of ξ and the covariance matrix are estimated directly from the data.

Chapter 3

Results

This chapter presents the results from applying the method described in Chapter 2 to the first- and fifth-year WMAP data releases. Even though the general methodology is the same, some of the technical details changed from one analysis to the next - these differences are pointed out in the relevant sections.

3.1 Year one

We use the peak-peak correlation function in a number of different ways to investigate the properties of the maps. We will use the following nomenclature: H for Hot, C for Cold, N for North and S for South:

- The most obvious way is to conduct a full-sky analysis in the unmasked regions of the maps, which we do for hot and cold spots separately - ξ_H and ξ_C.
- Motivated by a detection of a cold spot in the southern hemisphere and other hints of asymmetry (see Chapter 2), we also compute the peak-peak correlation function in each of the hemispheres individually, again looking at hot and cold spots separately in each case - ξ_{NH}, ξ_{NC}, ξ_{SH} and ξ_{SC}.
- In addition we look at the difference of correlation between the two hemispheres at a given angular scale and we define $\Delta\xi_H = \xi_{SH} - \xi_{NH}$ (similar for cold spots).
- Finally, we take the average of the peak-peak correlation function in the Northern

and Southern hemispheres in order to produce a computationally faster way to estimate the full-sky function - $\tilde{\xi}_H$ (similarly for cold spots).

For each we estimate ξ in 300 equally-spaced bins up to a maximum separation of 1800 arc minutes. Previously to computing equation (2.10) we rebin all data to 19 bins, of which we discard the first one[1]. Rebinning is necessary, otherwise C_{ij} is close to singular and numerically unstable to inversion. We explain each of our estimators in detail in the following sections.

3.1.1 All-sky analysis

We first consider all the hot spots above a certain threshold $\nu\sigma$ (or cold spots below a negative threshold $-\nu\sigma$) for the entire sky, except for the masked regions of galactic plane and point sources. The results for a threshold of $\nu = 1.5$ are shown in Figure 3.1. We also plot the peak-peak correlation function averaged over 100 Gaussian maps and the error bars on the Gaussian curve are the errors on the mean. The small error bars show good convergence of the average of the peak-peak correlation function from the 100 Gaussian maps. Figure 3.2 shows the convergence χ^2 for ξ_H and ξ_C with increasing number of maps. Although not optimally sampled, the structure we see at small angular scales is real structure, as expected from Heavens and Sheth (1999) and Heavens and Gupta (2001).

We see immediately that neither the hot spots nor the cold spots follow the Gaussian simulations - the cold spots show an excess of correlation whereas the hot spots show a lack of correlation with respect to the Gaussian simulations. These differences are, however, not significant; one disadvantage of the correlation function is that the errors can be highly correlated. The distribution of the χ^2 values for all of the Gaussian maps can be seen in Figure 3.1, together with the values for the WMAP data. We find both statistics are within the Gaussian 1σ confidence level. So the maps analysed in this way do not show any sign of non-Gaussianity. This is in agreement with Larson and Wandelt (2005) who also find no significant deviation from Gaussianity when they

[1]HEALPix defines neighbouring pixels as ones which share a pixel face. However, due to the highly variable pixel shapes in the surface of the sphere, these are not necessarily the closest pixels to the central one. This occasionally results in HEALpix selecting two very close pixels as being separate peaks which in turn results in unexpected (but explainable) features in the first few bins. Hence we choose to ignore these bins (which fall into the first one after rebinning). The effect these extra peaks have at large angular scales was tested for and found to be negligible.

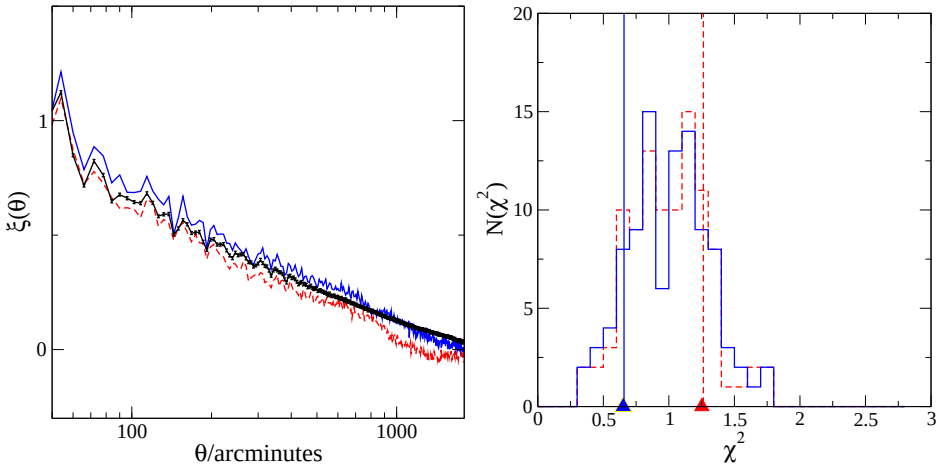

Figure 3.1: Left: the peak-peak correlation function of WMAP's data hot spots in the dashed (red) line and cold spots in the solid (blue) line. Simulated data (averaged over 100 Gaussian simulations) in the middle (black) line - the error bars shown are the errors on the mean. The threshold for peaks is $\nu = 1.5$. Right: The distribution of reduced χ^2 values for all of the 100 Gaussian maps: hot spots in the dashed (red) line and cold spots in the solid (blue) line. The χ^2 values for the WMAP data are represented by the small triangles and vertical lines

compute the peak-peak correlation of hot and cold spots in the first year data (although they work with lower resolution maps).

At the time this work was conducted, there were claims in the literature (see 2.1) concerning a cold spot in the southern hemisphere, and that the WMAP maps show an asymmetry in their statistical properties between the Northern and the Southern hemispheres, so we turn to this next.

3.1.2 North-South analysis

To further investigate any discrepancy between the WMAP data and our Gaussian simulations we estimate the peak-peak correlation function in the Northern and Southern hemispheres separately.

Figure 3.3 shows the peak-peak correlation function of the WMAP data for cold and hot spots calculated in the Northern and Southern hemispheres. We find a difference between the correlation of cold spots in the different hemispheres. Again we use a χ^2 statistic for $\Delta\xi_C$ and $\Delta\xi_H$, with the mean and covariance matrix estimated from 250

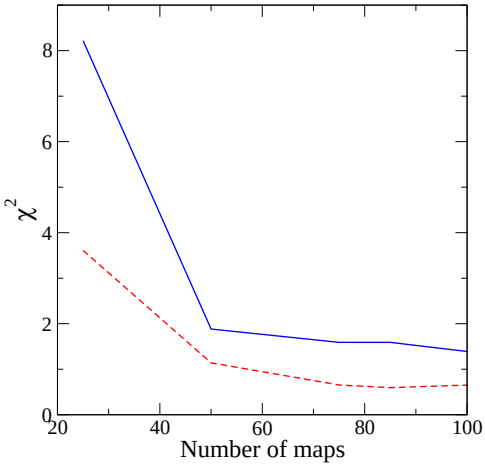

Figure 3.2: Convergence of ξ_C (solid, blue line) and ξ_H (dashed, red line) with number of Gaussian maps used to estimate $\bar{\xi}_i^G$ and C_{ij} as defined in Section 3.

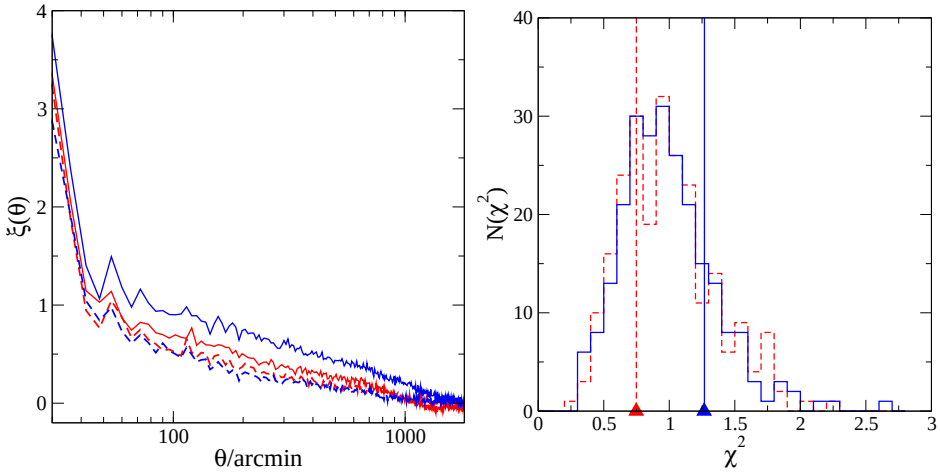

Figure 3.3: Left: the peak-peak correlation for the WMAP data in the two hemispheres - solid lines show the South and dashed lines the North. The inner pale (red) lines show hot spots and the outer (blue) lines show cold spots. The threshold for peaks is $\nu = 1.5$. Right: the distribution of reduced χ^2 for all 250 Gaussian maps. Hot spots in the dashed (red) line and cold spots in the solid (blue) line. The χ^2 values for the WMAP data are represented by the small triangles and vertical lines.

Gaussian maps. By analysing each hemisphere seperately, we are reducing the number of peaks available for the estimation of the peak-peak correlation function. Hence we found that a greater number of maps was needed to ensure good convergence of the average peak-peak correlation function and of the covariance matrix. See Figure 3.7 in the next section for convergence of some of the statistics with increasing number of maps.

We calculate χ^2 for our ensemble of Gaussian maps, whose distribution can be seen in Figure 3.3, together with the χ^2 value calculated for the WMAP data, for hot and cold spots.

We note that the fact we are finding the South-North difference not to be significant may be due to the fact that the peak-peak correlation function of threshold-selected peaks is highly sensitive to cosmic variance in the low multipoles. All the estimators are highly correlated and are shifted up and down in synchrony from Gaussian realisation to Gaussian realisation: the noisy low-ℓ multipoles can shift large numbers of peaks above or below the threshold depending on the mode amplitude. This suggests that the use of a high-pass filter - effectively removing the signal from cosmic variance for $\ell \leq \ell_{cut}$ - may be an efficient way to increase the sensitivity to non-Gaussian features.

3.1.3 Constraining in harmonic space

We construct an ensemble of 250 Gaussian maps, as described in section 2.3.3 for $\ell_{cut} = 0, 5, 10, 15, 20, 25, 30$ and 40. We compute $\Delta\xi_H$, $\Delta\xi_C$, ξ_{NH}, ξ_{NC}, ξ_{SH}, ξ_{SC}, $\tilde{\xi}_H$ and $\tilde{\xi}_C$ for all our Gaussian maps as well as the WMAP data. Figure 3.7 shows convergence of ξ_{SC} and $\Delta\xi_C$ with number of maps in the solid and dashed lines respectively. The same plot also shows the convergence of the same statistics but this time calculated in a single-frequency Q-band map (see section 4.5 for a single-frequency analysis).

Figure 3.4 shows $\Delta\xi_C(\theta)$ for some different ℓ_{cut} in the WMAP data. We note that the difference between the Southern and Northern hemispheres decreases as we remove more and more of the low order multipoles. This could be either due to the fact that cosmic variance alone is to blame for the North/South difference we see, or it could be due to the fact that whatever is causing this North/South difference is intrinsically a large scale effect.

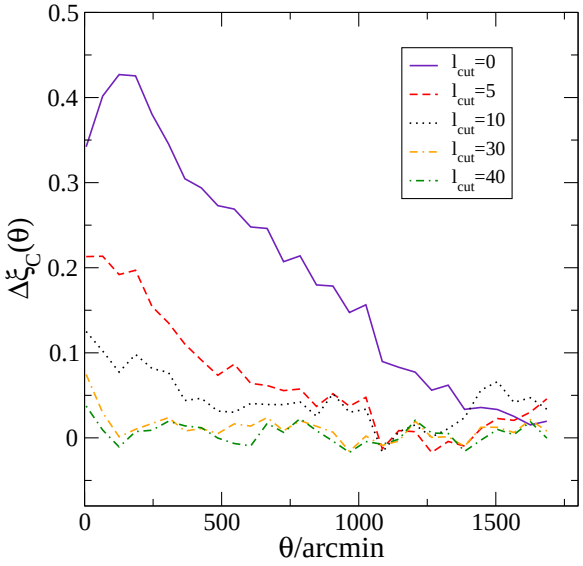

Figure 3.4: $\Delta\xi_C(\theta)$ for the WMAP map high-pass filtered with different values of ℓ_{cut}.

We test the significance of each of these differences by using χ^2. Figure 3.5 shows $\chi^2_{NS}(\ell_{cut})$ for cold and hot spots. We plot the distribution of χ^2 using *all* the different ℓ_{cut} Gaussian maps - these maps are not strictly independent (although the statistics share the same underlying χ^2 distribution over all values of ℓ_{cut}) so we use only the 250 independent maps at each ℓ_{cut} to draw conclusions about the significance of each detection - see section 3.1.7. The added histogram over the 2000 maps can be seen in Figure 3.5.

We do the same test and construct identical plots for all our statistics: ($\tilde{\xi}_H$,$\tilde{\xi}_C$) in Figure 3.6 and (ξ_{NH},ξ_{NC},ξ_{SH},ξ_{SC}) in the right panel of the same figure. The added histograms across all values of ℓ_{cut} for these statistics are very similar to that shown in Figure 3.5.

The first point to make is that the non-Gaussianity is consistently absent at $\ell_{cut} = 40$: there is no evidence from the peak-peak correlation function of non-Gaussianity on scales with $\ell > 40$.

The most significant non-Gaussian detections come from the cold spots in the Southern hemisphere, ξ_{SC}, at $\ell_{cut} = 10$, where we also find significant detections in

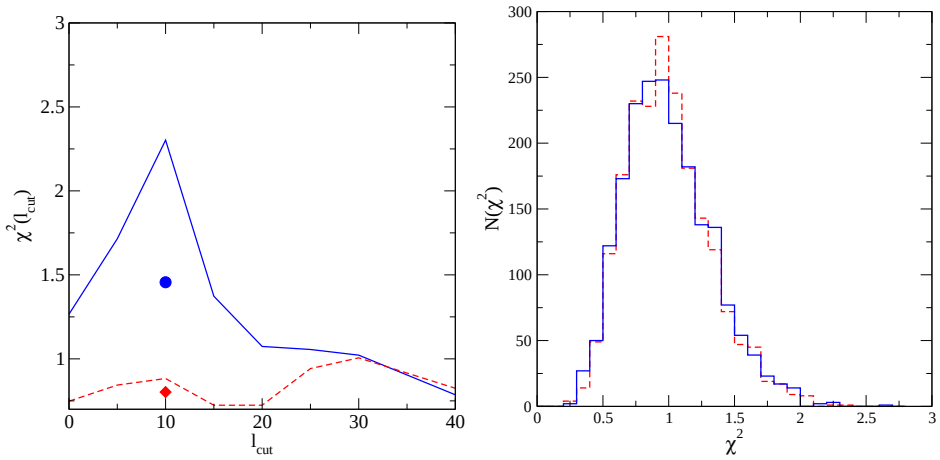

Figure 3.5: Left: χ^2 as a function of ℓ_{cut} for the WMAP data. Hot spots in the solid (red) line, cold spots in the dashed line (blue). The circle (blue) and the diamond (red) are the χ^2 value (cold spots and hot spots respectively) for runs with the regions of sky within 30 degrees of the galactic plane removed (see Section 3.1.4). Right: The added distribution of χ^2_{NS} values for all our Gaussian maps at all different ℓ_{cut}. Hot spots in the dashed (red) line, cold spots in the solid (blue) line. Similar histograms were produced for all of our other statistics, and all show a very similar added distribution of reduced χ^2 values.

the South-North difference for cold spots, $\Delta\xi_C$, and in the average of Northern and Southern hemispheres for cold spots, $\tilde{\xi}_C$. In addition to this, we have less significant detections at $\ell_{cut} = 20, 25$ and 30 in ξ_{SC} and $\tilde{\xi}_C$, see Figure 3.6. All of these do not appear in a North minus South analysis. This could be simply because the signal is not significant enough to show up in such analysis (we are roughly doubling the variance of our estimator by subtracting the data of the Sourthern and Northern hemispheres).

3.1.4 Constraining in real space

We further investigate the origin of this detection by removing extra regions near the masked Galactic plane. We work on the maps where the significance of the signal is the strongest (those with $\ell_{cut} = 10$), which we mask with an extended mask which additionally excludes all sky within 30 degrees of the galactic plane.

We proceed the same way as before and compute the full set of estimators: $\Delta\xi_H$, $\Delta\xi_C$, ξ_{NH}, ξ_{NC}, ξ_{SH}, ξ_{SC}, $\tilde{\xi}_H$ and $\tilde{\xi}_C$ for all our Gaussian maps as well as the WMAP

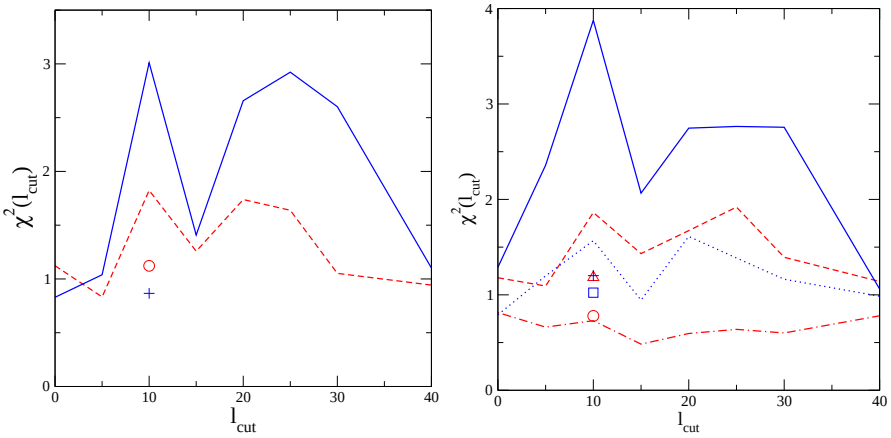

Figure 3.6: Left: $\chi^2(\ell_{cut})$ for $\tilde{\xi}_C$ solid (blue) line and $\tilde{\xi}_H$ dashed (red) line for the WMAP data. The single points at $\ell_{cut} = 10$ are the χ^2 values for cold spots (blue cross) and hot spots (red circle) in runs with the regions of sky within 30 degrees of the Galactic plane removed (see section 3.1.4). Right: $\chi^2(\ell_{cut})$ for ξ_{NC} (blue dotted line), ξ_{NH} (red dot-dashed line), ξ_{SC} (blue solid line) and ξ_{SH} (red dashed line) for the WMAP data. The points at $\ell_{cut} = 10$ are the χ^2 values for runs with the regions of sky within 30 degrees of the Galactic plane removed (see section 3.1.4): ξ_{NH} in the red circle, ξ_{NC} in the blue square, ξ_{SH} in the red triangle and ξ_{SC} in the blue cross.

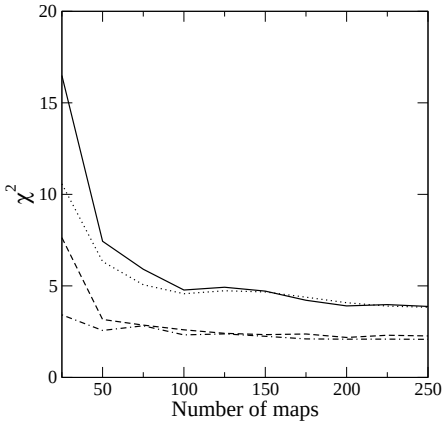

Figure 3.7: Convergence of some of our statistics which yielded detections of non-Gaussianity with increasing number of Gaussian maps used to estimate the mean and the covariance matrices. For the QVW map we show ξ_{SC} in the solid line and $\Delta\xi_C$ in the dashed line. For the single-frequency Q-band map we show ξ_{SC} in the dotted line and $\Delta\xi_C$ in the dot-dashed line

data and use the adequate χ^2 statistic for each of them to test the WMAP data for non-Gaussianity (we generate new random catalogues whose spatial distribution follows that of the new masks).

Figures 3.5 and 3.6 show how the new χ^2 values compare with the ones previously obtained when we did not use any extra galactic cut - all values drop significantly to values which are perfectly consistent with the Gaussian hypothesis (the most extreme value being for $\Delta\xi_C$), indicating that our significant non-Gaussian detection in the cold spots is located within 30 degrees of the galactic plane. This hints at residual foreground contamination associated with the Milky Way.

We note that we have only tested this on maps with $\ell_{cut} = 10$ since this is where we have found our strongest detection. We cannot discard the possibility that the effect that yields detections on maps with $\ell_{cut} = 15, 25$ and 30 is a different effect altogether which does not lie in the galactic region.

3.1.5 A single-frequency analysis

To check whether the non-Gaussian signal we detect is related to possible residual foregrounds in the WMAP data we conduct a single frequency analysis of the maps. Indeed, the expected Galactic foreground contribution to the WMAP maps consists mainly of synchrotron, free-free and dust emission. All of these effects are frequency-dependent and obviously non-Gaussian. If any foreground residuals are still present in the foreground-cleaned data then we would expect them to contribute differently to each of the different frequency maps. We note that any residual noise may also contribute differently to each frequency.

We construct the real map and each of the 250 simulated single frequency maps, at the Q, V and W bands. We then smooth the WMAP and Gaussian maps with a 12 arc minute FWHM Gaussian beam and high-pass filter with a $\ell_{cut} = 10$ window function (where we had the most significant non-Gaussian detection).

We calculate the full set of estimators for each of the frequencies: $\Delta\xi_H$, $\Delta\xi_C$, ξ_{NH}, ξ_{NC}, ξ_{SH}, ξ_{SC}, $\tilde{\xi}_H$ and $\tilde{\xi}_C$, for which the χ^2 values can be seen in Figure 3.8.

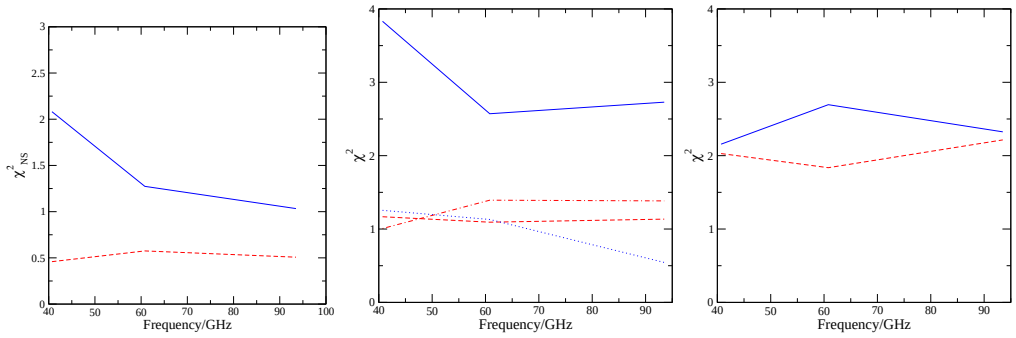

Figure 3.8: χ^2 for all three frequencies: Q (41 GHz), V (61 GHz) and W (94 GHz) on maps with $\ell_{cut} = 10$. Statistics for cold spots in the solid (blue) line, for hot spots in the dashed (red) line. Right: $\Delta\xi_H$ (red) and $\Delta\xi_C$ (blue). Middle: ξ_{NH} in the dot dashed (red) line, ξ_{NC} in the dotted (blue) line, ξ_{SH} in the dashed (red) line and finally for ξ_{SC} in the solid (blue) line. Right: $\tilde{\xi}_H$ in the dashed (red) line and for $\tilde{\xi}_C$ in the solid (blue) line.

We find significant non-Gaussian signals coming from the cold spots $\Delta\xi_C$ in the Q band and ξ_{SC} in all three bands, although it is strongest in the Q band. We also find detections in our full-sky estimates in the cold spots in all three bands, and, for the first time, in the hot spots in bands Q and W (left panel. Figure 3.8).

We may be seeing a frequency-dependent type of non-Gaussianity, although we can not put aside the possibility of a cosmological origin. To improve readability we do not present the plots with the χ^2 distributions of the 250 Gaussian maps for each of the frequencies and for each of the estimators. We do, however, quote the number of Gaussian maps with a $\chi^2_{Gaussian} \geq \chi^2_{WMAP}$ for all significant detections in Table 3.1, section 3.1.7.

3.1.6 Removing the cosmological signal

In order to investigate the possibility of any contributions from foregrounds or unexplained noise properties, we remove what is taken to be the cosmological signal from our analysis. To do so we subtract different single-frequency maps to produce three maps which contain only a mix of subtracted residual foregrounds (if any) and noise. We produce a $V - Q$, a $V - W$ and a $Q - W$ map, which are simply a pixel-by-pixel subtraction of each of the single frequency maps, constructed as described in section 2.2.4.

With the cosmological signal removed, the detailed noise properties of these three subtracted maps at large angular scales now become important for our analysis and one should be careful when constructing equivalent Gaussian maps (see 2.2.2). We therefore take a slightly different route to construct the Gaussian simulations with which we compare the WMAP data, and we now make use of the 110 noise simulations supplied by the WMAP team. We construct single-frequency noise maps by adding the respective individual radiometer simulations following the same weighting scheme as described in section 2.2.4, which we then smooth and high-pass filter with a $\ell_{cut} = 10$ window. We then subtract different frequency noise maps in order to produce 110 simulations with which we compare our real $V-Q$, $V-W$ and $Q-W$ maps. We re-emphasize that for maps which include the signal, the non-white nature of the noise at low-ℓ is essentially irrelevant, as the signal dominates entirely (Section 2.2.2).

We construct $\Delta\xi_H$, $\Delta\xi_C$, ξ_{NH}, ξ_{NC}, ξ_{SH}, ξ_{SC}, $\tilde{\xi}_H$ and $\tilde{\xi}_C$ for the simulations and the real data as before and use the respective χ^2 statistic to probe for non-Gaussian signatures. In this case, our total number of maps was constrained by the number of noise simulations provided by the WMAP team. Figure 3.9 shows how the reduced χ^2 values for $\Delta\xi_H$, $\Delta\xi_C$, ξ_{NH}, ξ_{NC}, ξ_{SH} and ξ_{SC} in the $Q-W$ map change with number of simulated maps used (the $Q-V$ and $V-W$ maps produced very similar results). The results show clear convergence to some value well within the 1σ confidence levels. The reason why we observe faster convergence in these maps could simply be due to the fact that we are removing the cosmological signal from the analysis and with it much of the variance.

Figure 3.9 also shows ξ_{NC} and ξ_{SC} for the WMAP data and also $\bar{\xi}^G_{NC}$ and $\bar{\xi}^G_{SC}$ where the average is done over the 110 simulated $V-Q$ noise maps.

Some comments on this figure are appropriate. Firstly we note that there is a large intrinsic North/South asymmetry in the Gaussian noise maps. This is due non-stationary noise due to the uneven scanning pattern of the WMAP satellite. We recall that pixel noise is weighted according to the number of times a pixel has been observed, and as such this feature is fully simulated in all our previous maps. This large-scale structure combined with the fact we are applying an asymmetric mask to the data results in the non-zero and North/South asymmetric peak-peak correlation function we

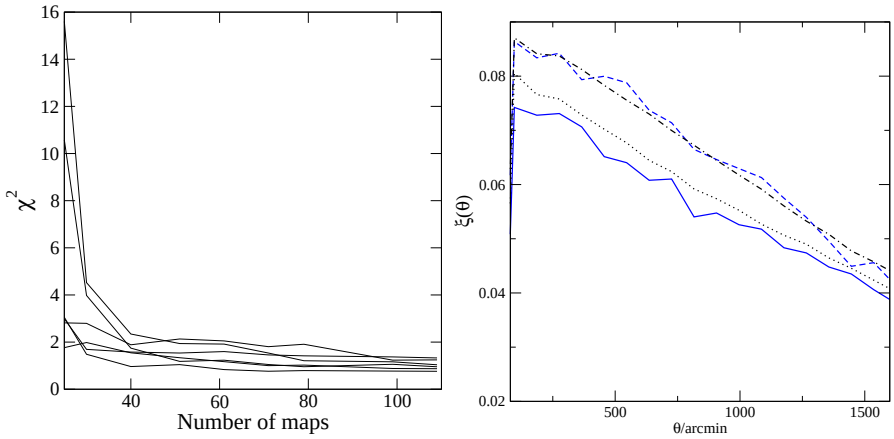

Figure 3.9: Left: Convergence of the reduced χ^2 values for $\Delta\xi_H$, $\Delta\xi_C$, ξ_{NH}, ξ_{NC}, ξ_{SH} and ξ_{SC} in the $Q-W$ map as a function of number of Gaussian maps. Right: ξ_{NC} and ξ_{SC} estimated from the $V-Q$ subtracted maps. WMAP's data are the solid line for Southern hemisphere and dashed line for Northern hemisphere (both in blue). Gaussian averaged data in dotted line for Southern hemisphere, dot dashed line for Northern hemisphere (both in black)

see. We draw attention to the fact that this asymmetry is qualitatively different from what we found in sections 3.1.2, 3.1.3, 3.1.4 and 3.1.5, since we now find an excess in correlation in the Northern hemisphere, as opposed to in the South[2]. This excess in correlation in the North is indeed seen in the Gaussian-averaged peak-peak correlation function of all our previous maps, although on a much smaller scale. Finally we note that there is a more noticeable deviation of the WMAP data from the Gaussian simulations in the Southern hemisphere. However, we find none of these to be significant. In fact, this statement extends to the other two cases: $V-W$ and $Q-W$. We find no signs of non-Gaussianity in any of the estimators in any of our combined noise and foreground maps, with all the χ^2 values well within values which are consistent with the Gaussian hypothesis (our most extreme χ^2 value comes from ξ_{SH} in the $V-Q$ map, where we find $\chi^2 = 1.49$ - see Table 3.1 in the next section for a summary of the most extreme values in all three maps).

[2]As a sanity test, we have also performed an identical analysis on purely white noise maps which include the WMAP's satellite scanning pattern and found them to have the same North/South asymmetry behaviour.

3.1.7 Summary

In this subsection we take the opportunity to summarise our results into one table and to elaborate on the confidence levels we have quoted throughout the paper. We do this by presenting a table with all the statistics for which we have found the WMAP1 data to have a reduced $\chi^2_{WMAP} \geq 2$, Table 3.1.

We recall that in all cases we have rebinned the data into 19 linearly-spaced bins, of which we use the last 18 to compute each of the χ^2 statistics. The P_{theory} column gives the probability of randomly obtaining a given value of $\chi^2 \geq \chi^2_{WMAP}$ assuming the underlying distribution is a χ^2 distribution with 18 degrees of freedom, and the $N_{Gaussian}$ column shows how many Gaussian maps have a $\chi^2 \geq \chi^2_{WMAP}$ for the corresponding estimator (the number in brackets in the total number of Gaussian maps). It is worth noting that the χ^2 distribution we estimate from the Gaussian maps fits a χ^2 distribution with 18 degrees of freedom which has been shifted slightly by $\Delta\chi^2 \approx 0.1$ to lower values. Hence any limit on high values of χ^2 based on this theoretical distribution is a conservative one. Shifting the Gaussian χ^2 distribution by $\Delta\chi^2 = 0.1$ results in the P_{theory} values in Table 3.1 roughly being halved.

We draw attention to our most striking detections, which come from the cold spots in the Southern hemisphere, appearing both in the co-added QVW map and in the single frequency Q band map with reduced χ^2 values of 3.877 and 3.831 respectively.

3.1.8 Conclusions on the first year analysis

Our main results are summarised in Table 3.1 in Section 3.1.7 - we find strong evidence for non-Gaussianity, mainly associated with the cold spots and with the Southern hemisphere; this non-Gaussianity disappears completely if we filter out the harmonic modes $\ell \leq 40$ and at least partially if we exclude sky within $|b| < 30^\circ$, so it is a large-scale effect associated with the galactic plane.

Recently, Larson and Wandelt (2005) have also used the peak-peak correlation function of cold and hot spots in their search for non-Gaussianity. Direct comparison of results is not straightforward as the resolutions of the maps used in the two studies are significantly different. However, in the simplest case where both groups looked at the full sky CMB temperature field (with equivalent masks based on the standard kp0 mask

Table 3.1: Our main detections. We present all situations which yielded vales of $\chi^2 \geq 2$. In addition to this and for the sake of completeness we also present the most extreme χ^2 values obtained in Section 3.1.6. P_{theory} is the theoretical probability of randomly obtaining a reduced $\chi^2 \geq \chi^2_{WMAP}$ assuming a reduced χ^2 distribution with 18 degrees of freedom and $N_{Gaussian}$ is the total number of Gaussian maps with $\chi^2 \geq \chi^2_{WMAP}$. In brackets is the number of Gaussian realisations used for each statistic.

Map	ℓ_{cut}	Estimator	χ^2_{WMAP}	P_{theory}	$N_{Gaussian}$
QVW	10	$\Delta\xi_C$	2.302	1.32×10^{-3}	0 (250)
QVW	5	ξ_{SC}	2.358	9.58×10^{-4}	0 (250)
QVW	10	ξ_{SC}	3.877	4.91×10^{-8}	0 (250)
QVW	20	ξ_{SC}	2.747	9.15×10^{-5}	0 (250)
QVW	25	ξ_{SC}	2.764	8.23×10^{-5}	0 (250)
QVW	30	ξ_{SC}	2.756	8.65×10^{-5}	0 (250)
QVW	10	$\tilde{\xi}_C$	3.011	1.71×10^{-5}	0 (250)
QVW	20	$\tilde{\xi}_C$	2.658	1.59×10^{-4}	0 (250)
QVW	25	$\tilde{\xi}_C$	2.923	3.01×10^{-5}	0 (250)
QVW	30	$\tilde{\xi}_C$	2.601	2.25×10^{-4}	0 (220)
Q	10	$\Delta\xi_C$	2.081	4.57×10^{-3}	2 (250)
Q	10	ξ_{SC}	3.831	6.78×10^{-8}	0 (250)
V	10	ξ_{SC}	2.571	2.70×10^{-4}	0 (250)
W	10	ξ_{SC}	2.729	1.02×10^{-4}	0 (250)
Q	10	$\tilde{\xi}_C$	2.156	3.02×10^{-3}	0 (250)
V	10	$\tilde{\xi}_C$	2.695	1.26×10^{-4}	0 (250)
W	10	$\tilde{\xi}_C$	2.325	1.16×10^{-3}	0 (250)
Q	10	$\tilde{\xi}_H$	2.029	6.04×10^{-3}	0 (250)
W	10	$\tilde{\xi}_H$	2.215	2.17×10^{-3}	1 (250)
V-Q	10	ξ_{SH}	1.494	8.10×10^{-2}	10 (110)
Q-W	10	$\Delta\xi_H$	1.328	1.58×10^{-1}	22 (110)
V-W	10	$\Delta\xi_C$	1.426	1.07×10^{-2}	10 (110)

applied), both results are in agreement in the sense that both fail to yield a detection. We believe this lack of detection is a result of large cosmic variance in low-ℓ multipoles.

We investigate this further by removing some of the low order multipoles from the maps, in the hope that by doing so we are increasing our sensitivity to non-Gaussian features by reducing the effects of cosmic variance. Once we remove all harmonic modes with $\ell \leq 10$ we systematically find anomalies related to the cold spots in the WMAP data and, when looking at both hemispheres separately, we not only find a striking North/South asymmetry, we repeatedly find the strongest anomalies to be in the Southern hemisphere. This is not unheard of: Vielva et al. (2004) first found an anomalous large cold spot in the Southern hemisphere (nicknamed *The Spot*), a detection which was followed by Cruz et al. (2005), Mukherjee and Wang (2004) and McEwen et al. (2005) and confirmed repeatedly. However, we do find that our detections disappear when we exclude sky regions within 30 degrees of the Galactic plane (we recall that *The Spot* is localised at approximately $(b = -57^\circ, l = 209^\circ)$, well outside our cut regions of sky). We therefore conclude that our detections come mainly from something other than *The Spot*.

We also find a difference between the northern and southern hemispheres. The asymmetry we find in this study seems to be a large scale effect, once again related only to the cold spots and to be contained within 30 degrees of the Galactic plane.

We investigate our detections further by firstly conducting an analysis in single frequency maps. We find some evidence for a dependence of the signal with frequency when we look at different hemispheres (peaking at 41GHz, corresponding to the Q band and in agreement with Liu and Zhang 2005), but this detection does not appear in a full-sky analysis. Secondly we remove the cosmological signal from the analysis by subtracting different frequency maps and testing the resulting foreground/noise combination maps for non-Gaussian signals. We find no signs of non-Gaussianity in these subtracted maps.

Finally we note that even though a contamination of residual point sources would affect the hot spots statistics, they would not show in the cold spots analysis.

How do we make sense of these results? A simple explanation seems untenable.

The fact that the signal becomes insignificant when the Galactic plane is removed suggests unsubtracted Galactic foregrounds are responsible; the large-scale nature of the signal is certainly consistent with this picture. One would then expect the individual frequency maps to show a significant signal, and this we do find, most strikingly in the Q band. However, the difference maps do not show a significant detection; these maps should directly test the residual foregrounds and noise, so the absence of detected non-Gaussianity does not obviously support this picture. We can reconcile these observations if the residual foregrounds affect more than one frequency band, and the subtraction removes the contamination to some extent. The fact that we find non-Gaussianity in all the single-frequency bands adds some support to this complex picture. In our view this is the most likely explanation for the results we find, but we cannot exclude a primordial origin for at least part of the non-Gaussian signal.

3.2 Year five

Given the unclear picture which emerged from our analysis of the first year WMAP data, it is interesting to revisit the problem with a new dataset. In this section we analyse WMAP's fifth-year data release and take the opportunity to extend our study by:

- considering the cross-correlation of peaks over different frequencies (equation 2.9),
- explicitly considering f_{NL} models, and
- considering the effect of the observed ISW effect in the peak-peak auto- and cross-correlation functions.

The differences in the WMAP data-analysis pipeline from year-one to year-five were mainly associated with a better estimation of the beam profile (section 2.2.2) and foreground contributions, both Galactic and extra-Galactic (section 2.2.1). As a consequence, and similarly to what happened for the 3rd year data analysis, the Q-band was excluded from the estimation of the temperature angular power spectrum in the 5th year data by the WMAP team. The reason is a combination of unremoved foregrounds and beam-asymetry problems (Hinshaw et al. 2007). This makes this band unsuitable for a non-Gaussianity analysis, not only due to the unremoved foregrounds but indeed also given the slight difference in the power-spectrum of the temperature fluctuations observed in the Q-band and of a Gaussian map, based on the published

best-fit C_ℓ and with the same instrumental properties. This potentially sheds some light over the results yielded by your first analysis of the data, as the Q-band stood out as being the least Gaussian of all frequencies analysed. We therefore choose to still analyse the CMB at this frequency in order to investigate whether our previous results will hold given the new dataset. However, we feel that any deviation from Gaussianity detected in this band alone should be attributed primarily to unsubtracted foregrounds.

We will use the following nomenclature: ξ_{QQ} corresponds to the auto-correlation function of peaks in the Q-band map (similar for other frequencies), ξ_{VW} to the cross-correlation function of peaks in the V- and W-band, and ξ_{V+W} to the auto-correlation of peaks in the $V + W$ co-added map.

Before any analysis, all maps are smoothed with a Gaussian beam of full-width half-maximum (FWMF) of 10 arc min. We estimate ξ in 49 equally spaced bins of 0.1 degrees, for separations between zero and five degrees. We increased the sampling of ξ and restrained our analysis to smaller scales, given that this is where most of the structure is located. We discard the first bins for the same reasons given in the footnote in section 3.1. We do not rebin the data - instead we use SVD to insure the inversion of the covariance matrix used is stable. In an attempt to beat down cosmic variance without having to remove the large-scale modes, we have reduced the threshold ν from $\nu = 1.5$ to $\nu = -1$. This has pros and cons, as discussed in section 2.3.5 and we will have a closer look at this issue later.

3.2.1 Full-sky analysis

The auto-correlation function

The simplest approach is to use all of the available sky and calculate the auto-correlation of peaks in the unmasked regions. We calculate ξ_{QQ}, ξ_{VV}, ξ_{WW} and ξ_{V+W} for 200 Gaussian maps with the respective noise and instrumental properties and for the observed CMB at the correspondent frequencies. Figure 3.10 shows ξ_{QQ} for hot and cold spots, along side with the mean estimated from the simulations (ξ_{VV}, ξ_{WW} and ξ_{V+W} look similar and are not plotted here). The error bars shown are calculated from the diagonal of the covariance matrix, and we can see that the variance from realisation to realisation is small, and the mean is well constrained. We have checked for convergence by looking at the evolution of χ^2 with increasing number of maps. This can be seen in Figure 3.11.

Figure 3.12 shows the distribution of reduced χ^2 values for these Gaussian realisa-

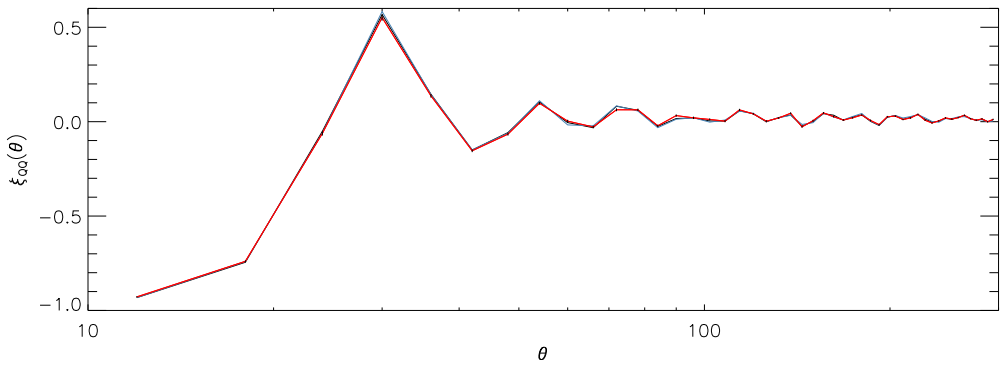

Figure 3.10: The auto-correlation function of hot (in the red) and cold (in the blue) spots in the Q-band, compared to the mean estimated from the Gaussian simulations. The error bars plotted on the line for hot spots are $C_{ii}^{1/2}$. θ is in arc-seconds.

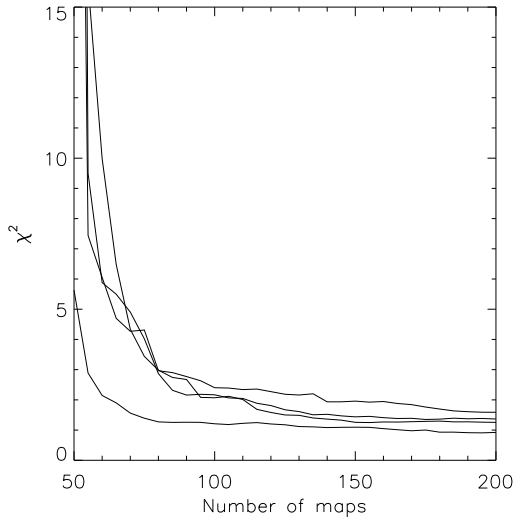

Figure 3.11: The evolution of χ^2 as we increase the number of maps for four of our estimators: $\xi_{QQ}, \xi_{VV}, \xi_{WW}$ and ξ_{V+W} for the cold spots. The curves for the hot spots are similar.

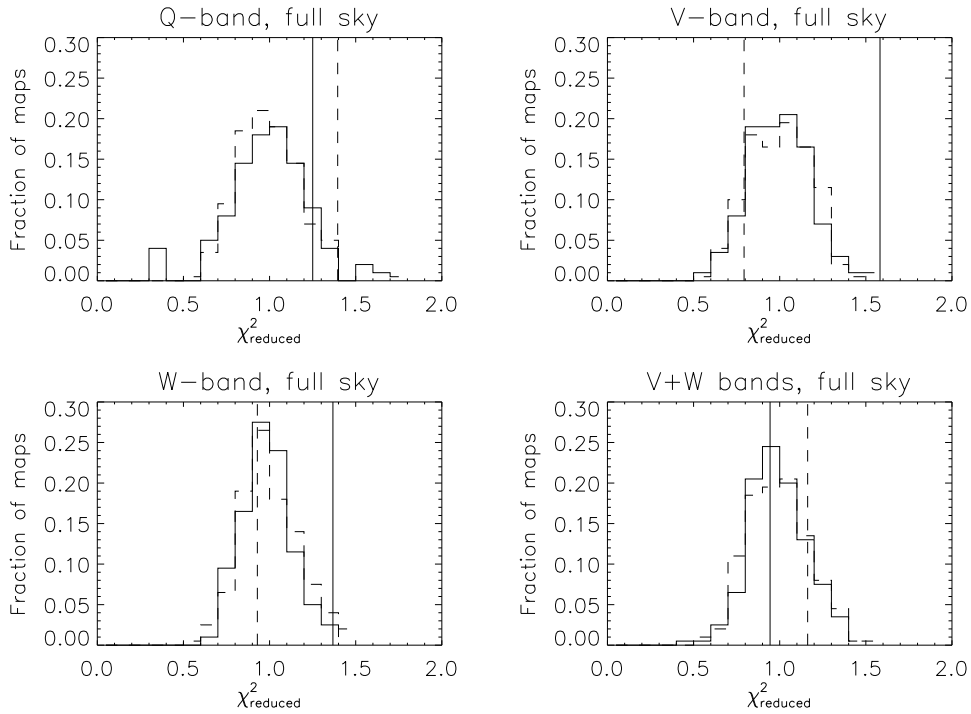

Figure 3.12: The distribution of values of reduced χ^2 for a set of 200 Gaussian simulated maps at the Q-, V- and W-frequencies, as well as a V+W co-added map. Coldspots are shown in the solid line, and hot spots in the dashed line. The χ^2 value of the observed CMB at each frequency is represented by the vertical line.

tions in each of the different frequencies. The value of the reduced χ^2 for the observed CMB is represented by a vertical line. We find that the data shows some signs of non-Gaussianity, in the V- and W-bands with cold spots, and in the Q-band with hot spots. As mentioned in section 2.2.4, the Q-band has a larger contribution from unremoved foregrounds than the other two bands, which likely explains the signal in the hot spots, but the signals in the V- and W-bands are worth investigating further.

The cross-correlation function

We also calculate the cross-correlation of peaks between the V- and W-frequencies, ξ_{VW}, for the same 200 Gaussian maps - see Figure 3.13. We see immediately that the cross-correlation function looks remarkably different from the auto-correlation function shown in Figure 3.10, particularly at small scales. The high power seen at small θ arises from the fact that even though the underlying temperature field is the same in both bands, differences in the beam profiles and noise properties mean that the same temperature peak generally falls on a different pixel in each of the frequencies.

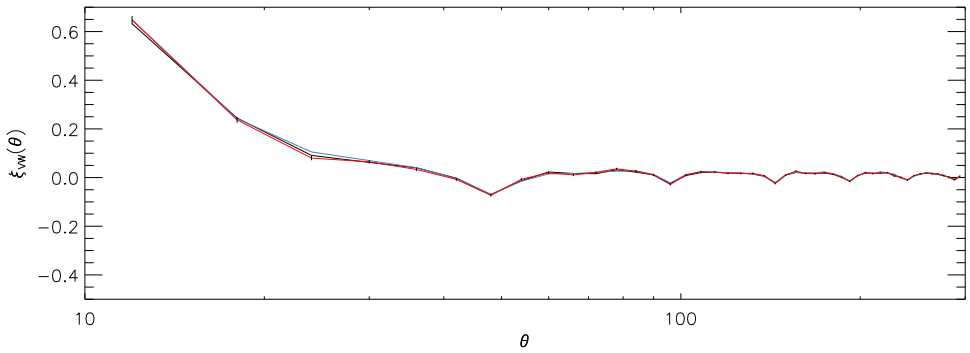

Figure 3.13: The cross-correlation function of hot (in the red) and cold (in the blue) spots between the V- and W-bands, compared to the mean estimated from the Gaussian simulations. The error bars plotted on the line for hot spots are $C_{ii}^{1/2}$. θ is in arc-seconds.

This means that there are peaks which are very close to each other across frequencies, although not at any given frequency.

Figure 3.14 shows the distribution of reduced χ^2 values for the Gaussian realisations for hot and cold spots. The goodness of fit value for the observed CMB is given by the vertical line - again we see some evidence for non-Gaussianity in the cold spots.

3.2.2 North-South analysis

The auto-correlation function

In the first-year data we detected a significant difference between the correlation of temperature peaks in the northern and southern hemispheres, relative to the galactic plane. As a follow up to that detection, we conduct a similar analysis on the year-five data. The procedure is identical to our full-sky analysis, but we extend the mask kq85 to exclude the northern or the southern hemisphere. We change our random catalogue accordingly. Similarly to what we found in our analysis of the first year data, we found that a higher number of maps was necessary to analyse each hemisphere separately. All of the north-south analysis in this and the next section were done using 300 maps - convergence can be seen in Figure 3.15.

The distribution of χ^2 values for each estimator is shown in Figure 3.16 and 3.17, for the northern and southern hemispheres. The V- and W-band signal associated with cold spots seen in the full-sky maps seems come mostly from the southern hemisphere. However, we see new detections when we look at each hemisphere separately which we

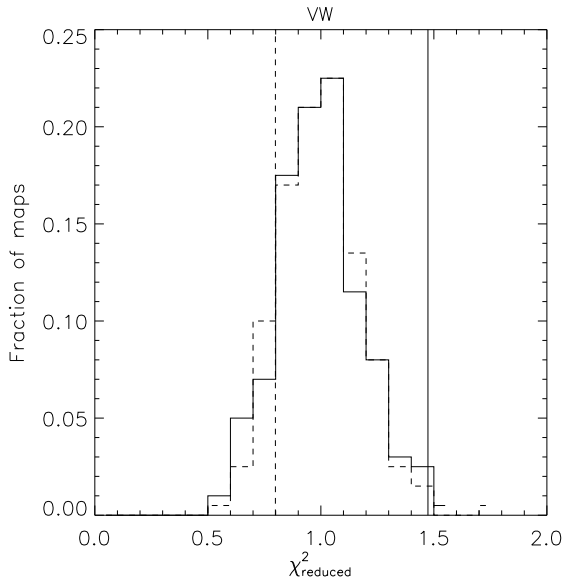

Figure 3.14: The distribution of values of reduced χ^2 for the cross-correlation of peaks between 200 Gaussian simulated maps at the V and W frequencies. Coldspots are shown in the solid line, and hot spots in the dashed line. The χ^2 values of the observed CMB are represented by the vertical line.

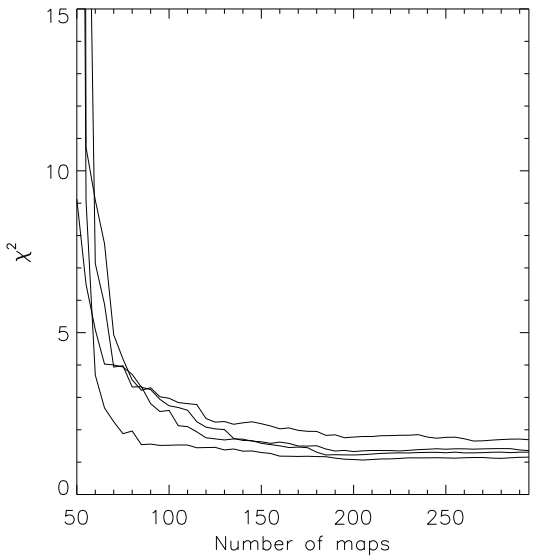

Figure 3.15: The evolution of χ^2 as we increase the number of maps for four of our estimators: $\xi_{QQ}, \xi_{VV}, \xi_{WW}$ and ξ_{V+W} for hot spots, in the northern hemisphere. Curves for the cold spots and the southern hemisphere are similar.

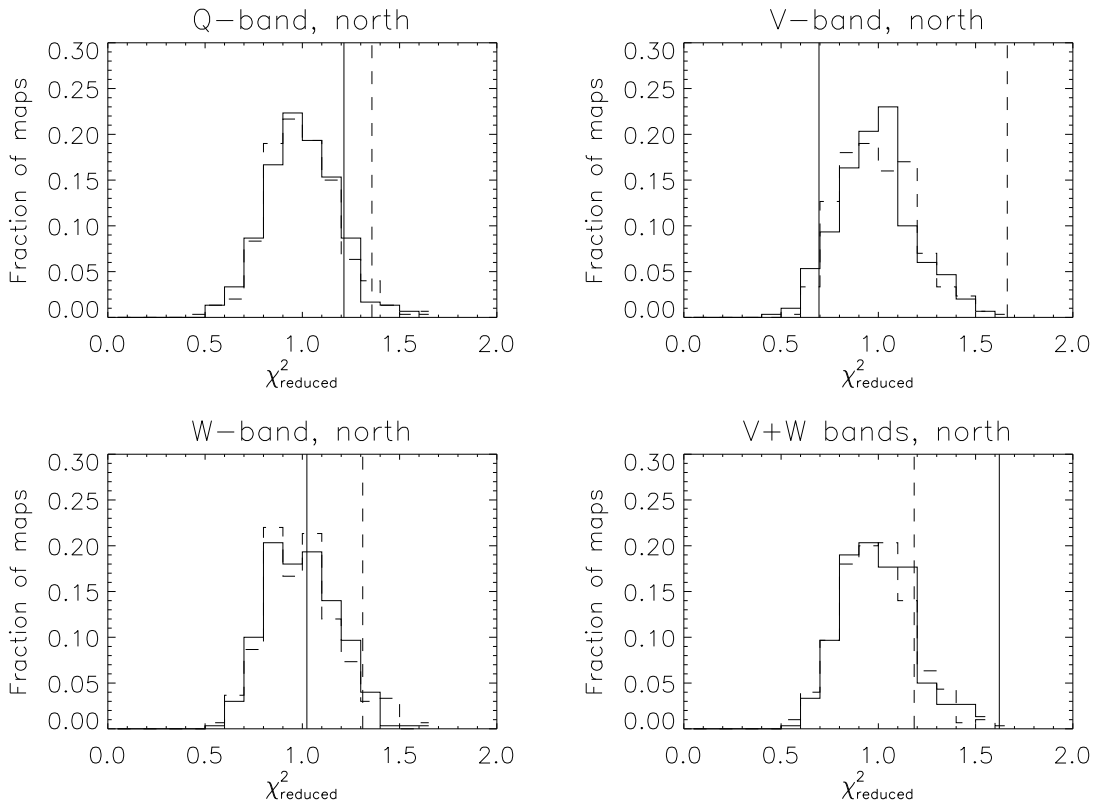

Figure 3.16: Histograms for the values of χ^2 for our auto-correlation estimators in the northern hemisphere. Coldspots are shown in the solid line, and hot spots in the dashed line. The χ^2 value of the observed CMB at each frequency is represented by the vertical line.

did not see in the full-sky analysis, especially associated with the hot spots, and most clearly in the V-band. A signal can appear in the two hemispheres separately, but not in the full-sky analysis, if they deviate from the mean with opposite signs. This is indeed what we see here.

The cross-correlation function

The cross-correlation between the V- and the W- bands in the two hemispheres can be seen in the top two panels of Figure 3.18. We see further evidence that the cold spots signal is predominantly coming from the south.

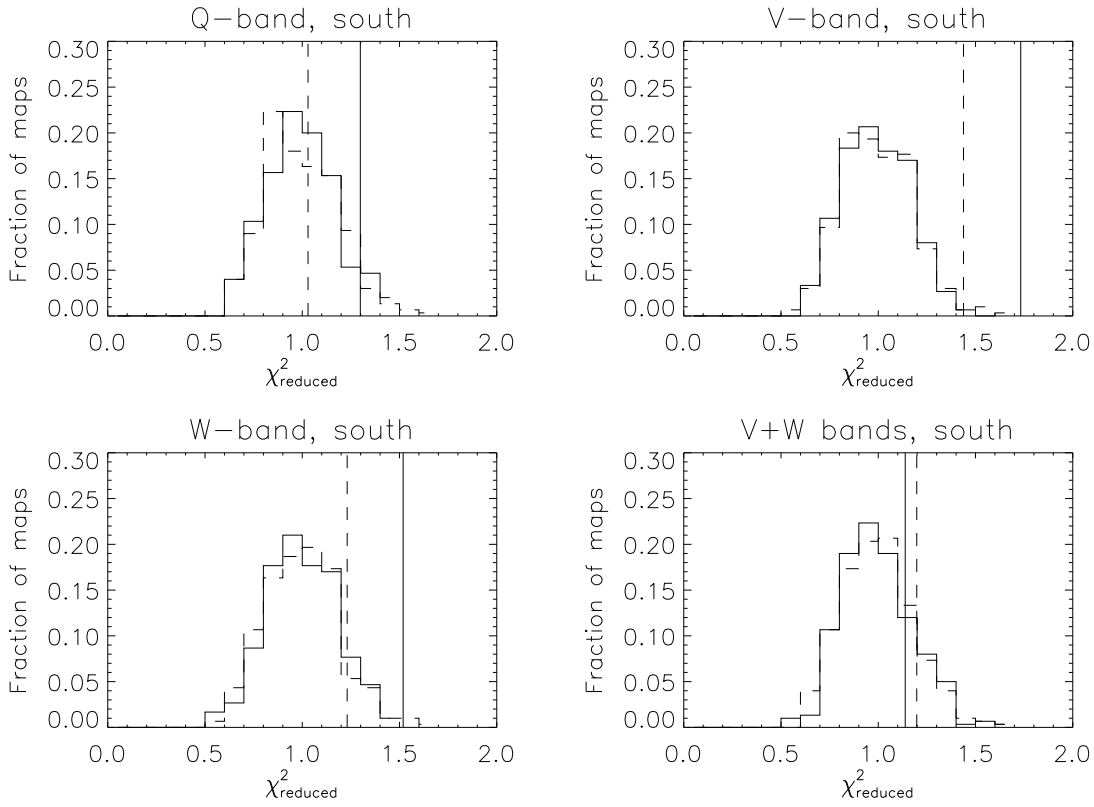

Figure 3.17: Histograms for the values of χ^2 for our auto-correlation estimators in the southern hemisphere. Coldspots are shown in the solid line, and hot spots in the dashed line. The χ^2 value of the observed CMB at each frequency is represented by the vertical line.

3.2.3 Constraining in real space

The auto-correlation function

Again prompted by what we found in our first-year analysis, we remove all regions around the Galactic plane for which $|b| < 30$ degrees. The resulting histograms can be seen in Figures 3.19 and 3.20. The values of χ^2 in the southern hemisphere are now all fully consistent with Gaussianity. The anomaly in the hot spots seen in the V-band remains although with a smaller significance level, and surprisingly we now see a signal in the cold spots which was unseen before, predominantly in the V-band but also in the W-band.

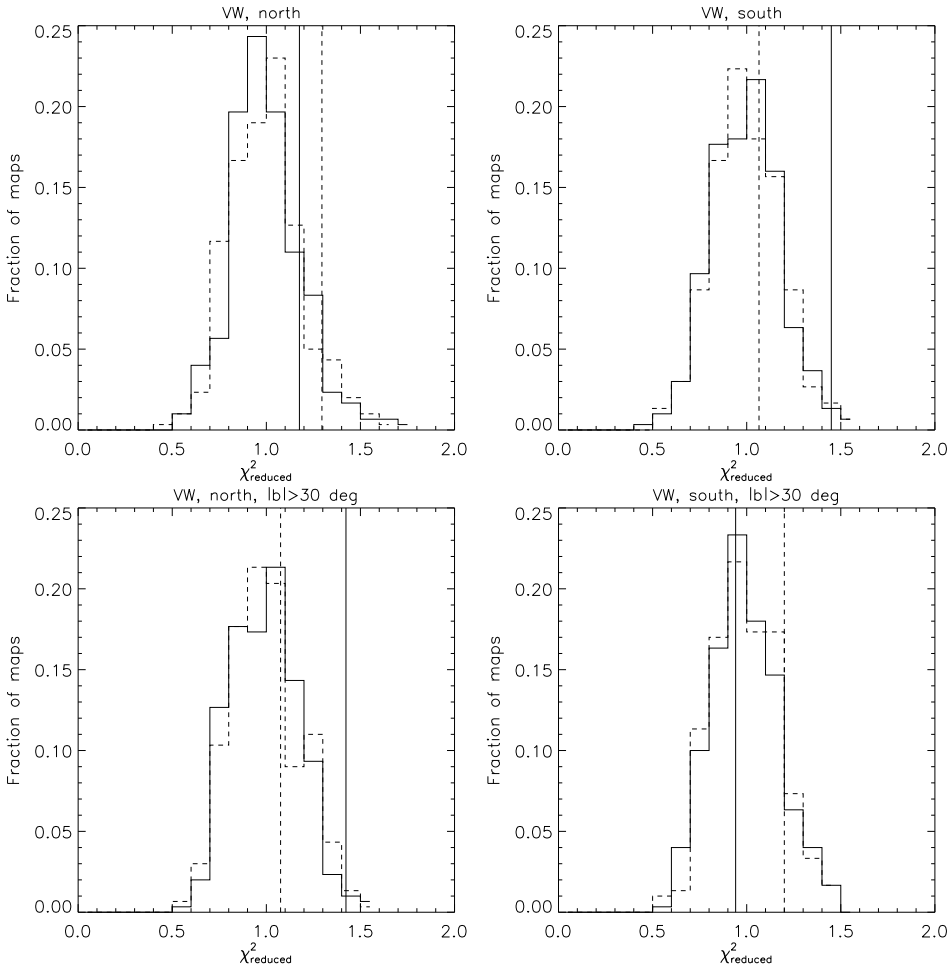

Figure 3.18: χ^2 distribution for ξ_{VW} in the northern and southern hemispheres. Top two panels show the results using the KQ85 mask (section 3.2.2), and the bottom two panels show the results excluding the regions within 30 degrees of the galactic plate (section 3.2.3).

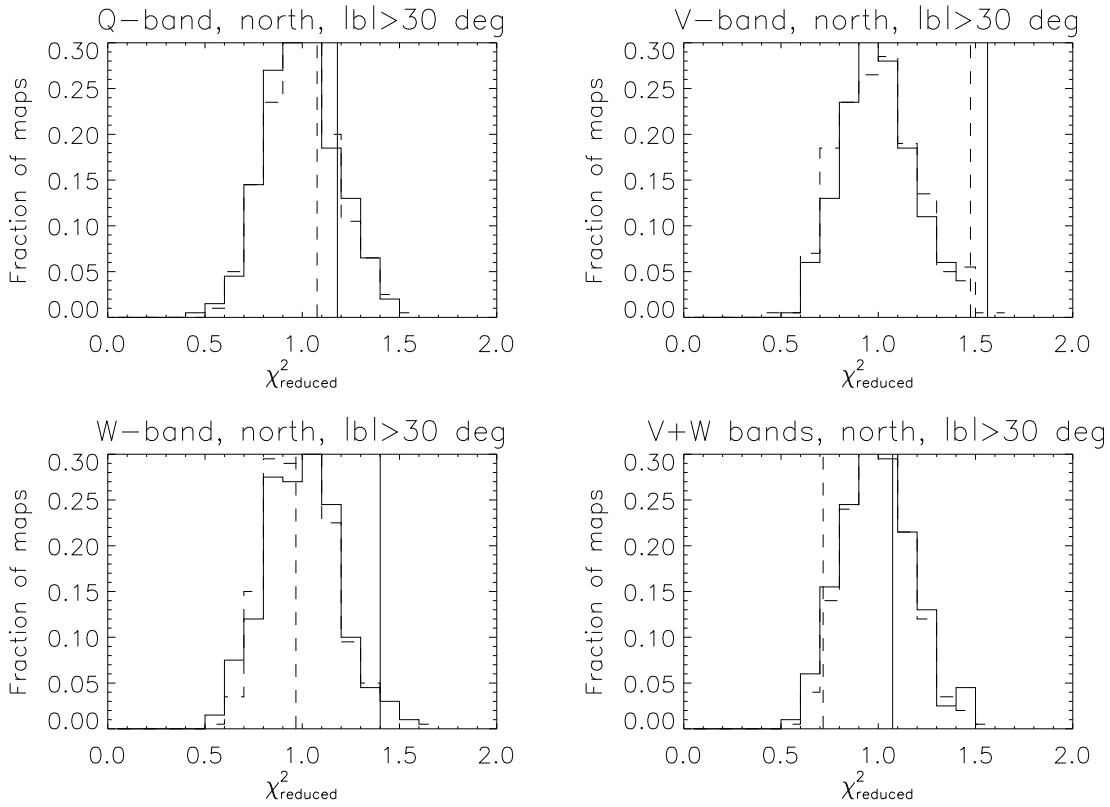

Figure 3.19: Histograms for the values of χ^2 for our auto-correlation estimators in the regions of sky out-width 30 degrees of the galactic plate, in the northern hemisphere. Coldspots are shown in the solid line, and hot spots in the dashed line. The χ^2 value of the observed CMB at each frequency is represented by the vertical line.

The cross-correlation function

The results for the cross-correlation function of temperatures out-with 30 degrees of the Galactic plane can be seen in the bottom two panels of Figure 3.18. We see the same behaviour as we saw with the auto-correlation function results: the signal in the southern hemisphere disappears, but we see a signal emerging in the northern hemisphere which we did not see before.

This curious signal, present both in the cross- and auto-correlation functions suggests a very unclear picture. The most immediate explanation is that we are seeing a localised source of non-Gaussianity which is too weak to show up when we analyse the whole northern hemisphere. If this is the case, we would expect the signal to come predom-

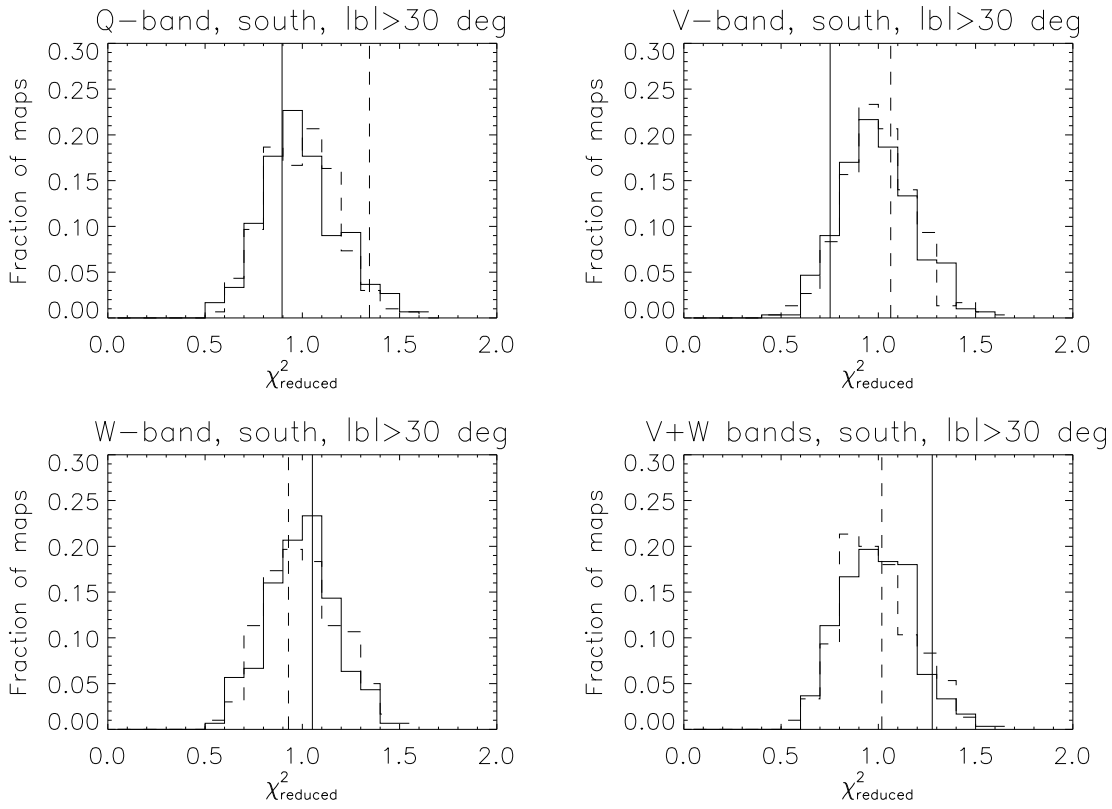

Figure 3.20: Histograms for the values of χ^2 for our auto-correlation estimators in the regions of sky out-width 30 degrees of the Galactic plate, in the southern hemisphere. Cold spots are shown in the solid line, and hot spots in the dashed line. The χ^2 value of the observed CMB at each frequency is represented by the vertical line.

inantly from scales associated to its angular size. We estimate the contribution to χ^2 from each scale by calculating

$$\chi^2_i = \sum_i \frac{(\xi_i - \bar{\xi}_i)^2}{C_{ii}} \tag{3.1}$$

which although ignores the correlation between non-adjacent scales, might provide insight about the cause of our results. Figure 3.21 shows χ^2_i for the estimators ξ_{VV} and ξ_{WW}, as calculated in the northern hemisphere for $|b| > 30$ degrees. We see that even though the signal comes from specific scales in each of these frequencies, it does not seem to be caused by the same scales in each frequency. This only adds to the difficulty of interpreting what is causing the appearance of this signal. The signal is robust to different bin widths, matrix-inversion methods and number of maps used. Its origin

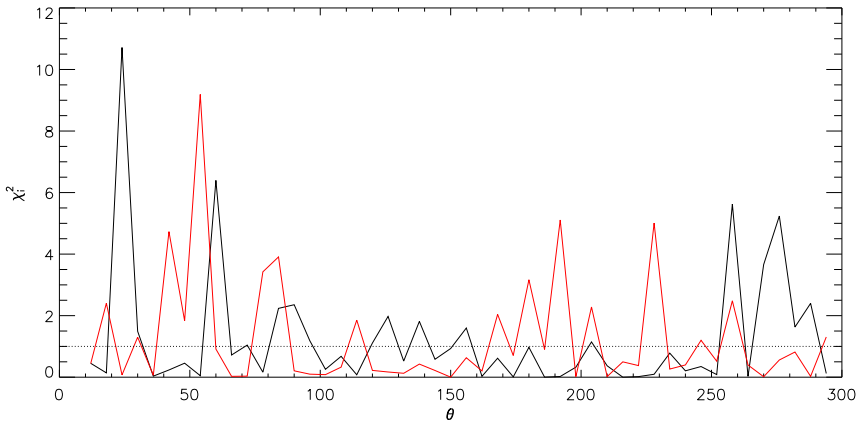

Figure 3.21: χ_i^2 for the estimators ξ_{VV} (black line) and ξ_{WW} (red line), as calculated in the northern hemisphere for $|b| > 30$ degrees. The dotted line shows the predicted mean value of one.

remains unknown to the time of writing.

3.2.4 The integrated Sachs-Wolfe effect

All of our detection of non-Gaussianity in the WMAP5 data, as they stand, offer very little insight about the causes behind them. Next we explore the possibility that a well known physical mechanism is behind at least part of our signal. The integrated Saches-Wolfe effect is a large-scale signal present the observed CMB which arises from the fact that CMB photons travel through evolving potential wells in their paths to us. The late ISW refers to changes occurring in the recent Universe. Francis & Peacock (2008, in prep) use the 2MASS survey in order to produce a reconstruction of the local density field, that together with a cosmological model - which describes the dynamics of the local Universe - can be used to calculate the contribution of the low-redshift density field to the late ISW contamination of the CMB. Using a local density field estimated up to $z = 0.3$, they have produced a late ISW temperature map, seen in Figure 3.22. In practice we analyse two reconstructed ISW maps, produced using two different methods but using the same dataset. In principle, the two maps represent the same thing and they should give identical results. Significant differences would indicate some error associated with the method for reconstruction. We will refer to these maps as the 2D and the 3D reconstructions (more details in Francis and Peacock 2008, in prep).

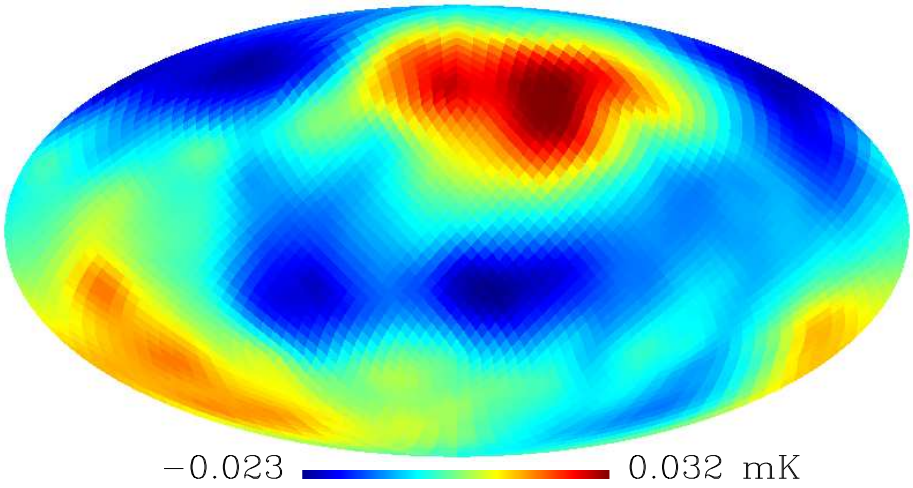

Figure 3.22: The ISW contribution to the observed CMB, as calculated by Francis & Peacock 2008 (in prep) using an estimation of the local density field up to z=0.3 and the cosmological model described in the text.

It is of interest to see how much the ISW signal affects different non-Gaussianity detections and other CMB anomalies. In this section we remove the ISW temperature map from the observed, foreground-reduced frequency maps. This gives us a CMB signal which is closer to the primordial CMB than the observed. We use the same Gaussian simulations as in section 2.3.2 given that the difference to the power-spectrum is small. We do the following

1. We subtract the predicted ISW map from the observed, full-resolution foreground-cleaned frequency maps.
2. We smooth the maps with a Gaussian beam of FWHM = 10 arc-min.
3. We re-mask the resulting map and remove the residual monopole and dipole.

In general, we find that the difference in the estimators induced from removing the estimated ISW signal from the maps is small, resulting in the goodness of fit values fluctuating with $\Delta\chi^2_{reduced} \approx 0.05$. The results are summarised in the Tables 3.2 and 3.3.

We only find one case in which a detection falls down to a level consistent with Gaussianity - ξ_{VV}, in hot spots and in the southern hemisphere sees its value of χ^2

reduced from 1.43 to 1.27 in the 2D reconstruction and 1.25 in the 3D reconstruction. However, given the fact that this is a lone event, we feel that it is unwise to over-interpret it. We conclude that the ISW does not play a role in the signals we see.

3.2.5 Summary

Table 3.4 summarises our detections in the 5th-year data. Our analysis have revealed a complicated picture which does not lend itself to any simple explanation. In spite of this, we can make the following statements:

1. All signals are frequency-dependent - we found no consistent signal over the three frequency bands in any of the cases.

2. Signals associated with the southern hemisphere (cold spots) disappear when we exclude the regions of sky within 30 degrees of the galactic plane.

3. Anomalies associated with hot spots are heavily associated with the northern hemisphere.

3.2.6 Conclusions on the fifth-year analysis

Once again we have found signs of non-Gaussianity in the WMAP data. We find that, by large, the signal that we found in the first-year data remains in the data today: once more we see an anomaly which is associated with cold spots, the southern hemisphere, and which disappears when we exclude the sky within 30 degrees of the galactic plane. We continue to see some asymmetry in data, with each hemisphere showing qualitatively different signatures of non-Gaussianity. Finally, we see a clear frequency dependence, with none of the maps analysed showing *consistent* departures from Gaussianity.

Map	Before ISW removal	After ISW removal (2D reconstruction)	$\Delta\chi^2$ (2D reconstruction)	After ISW removal (3D reconstruction)	$\Delta\chi^2$ (3D reconstruction)
Q, full-sky	1.3952871	1.4681376	-0.072850528	1.4846097	-0.089322581
V, full-sky	0.79310205	0.87312254	-0.080020498	0.88383181	-0.090729770
W, full-sky	0.92917691	0.93906113	-0.0098842192	0.96733246	-0.038155552
VW, full-sky	1.1623042	1.2101251	-0.047820921	1.1391697	0.023134443
Q, north	1.3573654	1.4127977	-0.055432266	1.4380926	-0.080727208
V, north	1.6602231	1.7218205	-0.061597424	1.6835322	-0.023309068
W, north	1.3099210	1.2732031	0.036717923	1.2816250	0.028295979
VW, north	1.1800490	1.1823953	-0.0023462838	1.1682028	0.011846198
Q, south	1.0284363	1.0471984	-0.018762142	1.0300560	-0.0016197646
V, south	1.4340021	1.2696905	0.16431159	1.2511920	0.18281010
W, south	1.2267701	1.2346055	-0.0078353685	1.2146827	0.012087417
VW, south	1.1967236	1.2438003	-0.047076727	1.0989857	0.097737822
Q, north, b30	1.0762889	1.2270067	-0.15071785	1.1857015	-0.10941267
V, north, b30	1.4738409	1.3970624	0.076778518	1.3752341	0.098606821
W, north, b30	0.96735869	1.0356369	-0.068278205	0.98193783	-0.014579136
VW, north, b30	0.71655327	0.70757045	0.0089828194	0.67361681	0.042936464
Q, south, b30	1.3449225	1.5053141	-0.16039165	1.3554369	-0.010514465
V, south, b30	1.0637318	1.1492181	-0.085486298	1.0901810	-0.026449264
W, south, b30	0.92875987	0.91198110	0.016778766	0.92085898	0.0079008842
VW, south, b30	1.0181098	1.1449669	-0.12685710	1.1424981	-0.12438834

Table 3.2: The change in the value of $\chi^2_{reduced}$ for our hot spots statistics estimators due to the removal of the local ISW effect, as shown in Figure 3.22. Results for cold spots shown in Table 3.3.

Map	Before ISW removal	After ISW removal (2D reconstruction)	$\Delta\chi^2$ (2D reconstruction)	After ISW removal (3D reconstruction)	$\Delta\chi^2$ (3D reconstruction)
Q, full-sky	1.2503281	1.2631188	-0.012790703	1.2882714	-0.037943363
V, full-sky	1.5823175	1.6095117	-0.027194188	1.6341854	-0.051867965
W, full-sky	1.3675620	1.2981338	0.069428233	1.2926912	0.074870805
VW, full-sky	0.94410902	0.89016865	0.053940366	0.93707981	0.0070292085
Q, north	1.2111371	1.1942278	0.016909282	1.2035917	0.0075453963
V, north	0.69247865	0.77365559	-0.081176943	0.80510400	-0.11262536
W, north	1.0209923	1.0615326	-0.040540314	1.0966578	-0.075665517
VW, north	1.6188697	1.5731120	0.045757675	1.6602142	-0.041344575
Q, south	1.2959187	1.3493766	-0.053457891	1.4198268	-0.12390808
V, south	1.7318566	1.7736560	-0.041799365	1.7548880	-0.023031425
W, south	1.5134875	1.5098637	0.0036238470	1.5001770	0.013310541
VW, south	1.1418154	1.1615992	-0.019783743	1.1540705	-0.012255101
Q, north, b30	1.1804753	1.1118426	0.068632766	1.1698382	0.010637094
V, north, b30	1.5618955	1.5143797	0.047515794	1.4868579	0.075037616
W, north, b30	1.4007526	1.4206091	-0.019856493	1.4868357	-0.086083103
VW, north, b30	1.0740224	1.0773703	-0.0033479143	1.0521836	0.021838821
Q, south, b30	0.89582892	0.98654480	-0.090715882	0.99464386	-0.098814943
V, south, b30	0.75207740	0.74092975	0.011147656	0.80429209	-0.052214688
W, south, b30	1.0515784	0.97421249	0.077365906	1.0115846	0.039993788
VW, south, b30	1.2770933	1.3112545	-0.034161187	1.3543291	-0.077235763

Table 3.3: The change in the value of $\chi^2_{reduced}$ for our cold spots statistics estimators due to the removal of the local ISW effect, as shown in Figure 3.22. Results for hot spots in Table 3.2.

In addition, we see some unexplained features in the fifth-year data, namely the appearance of a signal in the coldpots in the V-band (and the W-band to a lower significance) only when we mask out the regions within 30 degrees of the Galactic plane. These two signals are robust to numerical tests (such as number of maps or size of data vector), and the fact that they appear both in the auto- and cross-correlation functions suggests that they are caused by a real feature present in the data. The fact that we only detect them when we exclude the regions outwith the Galactic plane suggests a localised feature might be responsible but we could not associate the two signals with a single physical scale in the sky. At the time of writing we have been unable to explain what is behind these two detections.

The erratic nature of our detections suggests that we are seeing something which is not cosmological in origin. We have investigated whether the ISW could be the cause of part of the signal we see, but we have found this is not likely to be be the case.

As we emphasised before, whilst point sources might be behind the signal we see in the hot spots, they do not affect cold spots.

One possibility is that unsubtracted (or over subtracted) foregrounds remain in the data. A residual component with a frequency dependence would explain at least part of what we see. We finally conclude that this continues to be the most likely explanation to what we see.

3.3 f_{NL} constraints

We now turn to a different type of analysis, in which we assume an f_{NL} model and test the sensitivity of our estimators to changes in f_{NL}. This is in principle a much more rewarding approach to searching for non-Gaussianity, since we can directly infer a physical reason behind a detection.

At the time of writing this thesis we lack CMB simulated maps which incorporate both the correct cosmology from WMAP's 5th year results and have $f_{NL} \neq 0$. Our method requires the power-spectrum of the simulated Gaussian maps to match that of the observations, which cannot be done with the present maps. However, the cosmology is close enough to allow us to investigate how sensitive our estimators are to changes in f_{NL}, and we test this in maps of Planck resolution. With the correct maps it will

Map	Estimator	χ^2_{WMAP}	$N_{Gaussian}$
Q, full-sky	Hot spots	1.39	7 (200)
V, full-sky	Cold spots	1.58	0 (200)
W, full-sky	Cold spots	1.37	5 (200)
V&W, full-sky	Cold spots	1.47	1 (200)
Q, north	Hot spots	1.35	10 (300)
V, north	Hot spots	1.66	1 (300)
W, north	Hot spots	1.31	14 (300)
V&W, north	Cold spots	1.62	0 (300)
V, south	Hot spots	1.43	2 (300)
V, south	Cold spots	1.73	0 (300)
W, south	Cold spots	1.51	3 (300)
V&W, south	Cold spots	1.44	3 (300)
V, north, $\|b\| > 30$	Cold spots	1.56	1 (300)
V, north, $\|b\| > 30$	Hot spots	1.47	1 (300)
W, north, $\|b\| > 30$	Cold spots	1.40	6 (300)
V&W, north, $\|b\| > 30$	Cold spots	1.42	4 (300)

Table 3.4: The main indications of non-Gaussianity in the 5th-year data. $N_{Gaussian}$ is the total number of Gaussian maps with $\chi^2 \geq \chi^2_{WMAP}$. In brackets is the number of Gaussian realisations used for each statistic.

be straightforward to calculate the range of f_{NL} values allowed by the WMAP5 data.

We construct 200 maps as described in section 2.3.4, at the frequencies of 70 and 100GHz for $f_{NL} = [-100, 0, 30, 40, 50, 70, 100, 200]$. For each map we then calculate ξ_{70}, ξ_{100} and $\xi_{100,70}$ where the first two are the auto-correlation of peaks in the 70GHz and 100GHz respectively, and the third is the cross-correlation of peaks between the two frequencies. We do this for hot and cold spots, and then construct a data vector consisting of the two arrays: $y_i = \xi^H$ for $i = 1, \ldots, k$ and $y_i = \xi^C$ for $i = k+1, \ldots, 2k$, where H and C stand for hot and cold respectively and $k = 50$ is the number of bins in which each function is estimated. Whereas the peak-peak correlation function of hot and cold spots is the same in a Gaussian map, this is not true for $f_{NL} \neq 0$. Thus adding the two data vectors potentially increases the sensitivity of our estimators to changes in f_{NL}. Due to the large size of the data-vector we find the need to rebin the data to 25 data points for hot and cold spots, giving us 50 estimators in total. We

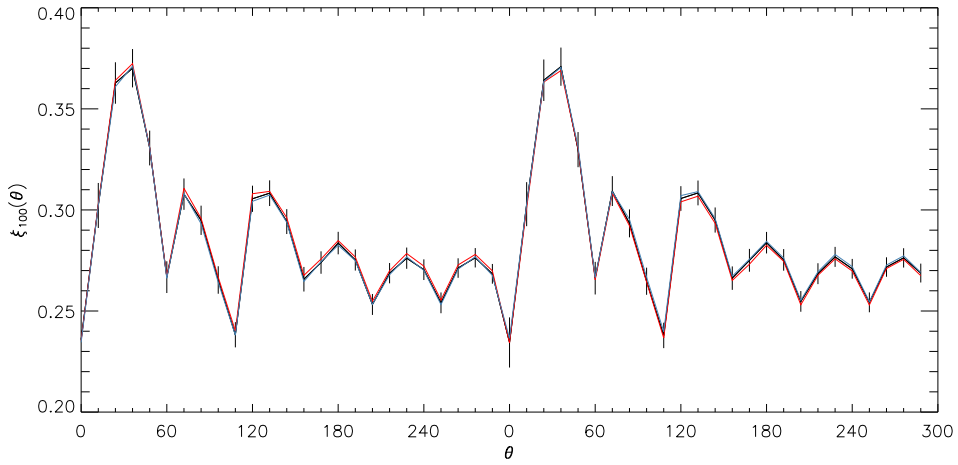

Figure 3.23: $\xi_{100-100}$ averaged over 200 maps for $f_{NL} = 40$ (black line), $f_{NL} = -100$ (blue line) and $f_{NL} = 200$ (red line). We show the curve for cold spots tagged at the end of the curve for hot spots - this is the data vector used in the analysis (see text). The error bars plotted on the black like are $C_{ii}^{1/2}$. The temperature threshold for the selection of peaks is $-\sigma$. Even though there is a qualitative change in the mean correlation function as a function of f_{NL} which is different for hot and cold spots, these changes are well within the 1σ level.

initially include all peaks above -1σ.

Figure 3.23 shows the average auto-correlation function in the 100GHz band, for $f_{NL} = 0, 40$ and 200, calculated over 200 maps. In the same figure we also plot $C_{ii}^{1/2}$ for the $f_{NL} = 40$ case, which shows clearly that the scatter from one realisation to the next is large compared to the differences in the models we are trying to differentiate.

This can also be seen by looking at the probability distribution of χ^2 given an f_{NL} model, for an assumed true value of f_{NL}. Figure 3.24 shows this for a test value of $f_{NL} = 40$. Even though there is a shift in the centre of these distributions, this illustrates by eye how distinguishing between these values of f_{NL} with only one observed CMB is simply too ambitions. We see similar results with the cross-correlation function.

3.3.1 Summary

Even though this first attempt suggests that the sensitivity to f_{NL} of the auto- and cross-correlation function of peaks is far from being competitive, there are routes still to be explored. One of them is to increase the threshold, for the reasons mentioned in section 2.3.6. Even though this undoubtedly increases the variance of our estimators

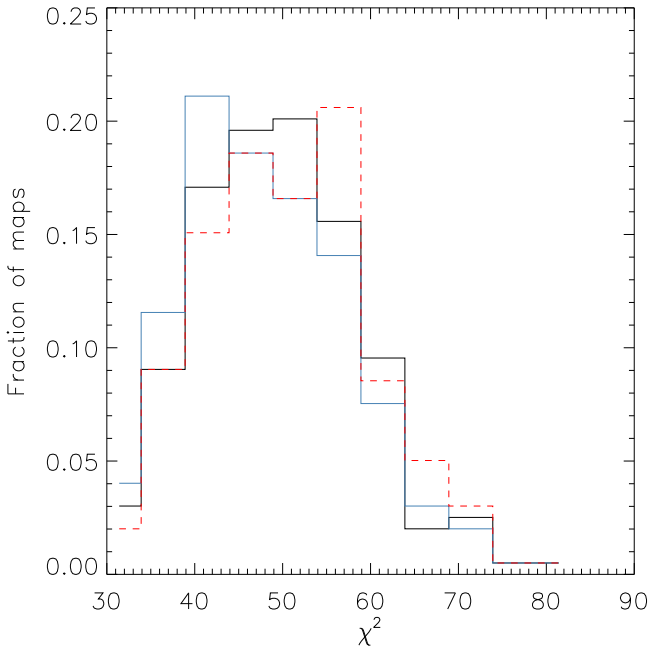

Figure 3.24: The probability distribution of χ^2 for three values of f_{NL} assuming a true value of $f_{NL} = 40$, as estimated from 200 maps in each case. $f_{NL} = 40$ in the thick black line, $f_{NL} = -100$ in the blue line, and $f_{NL} = 200$ in the red dashed line.

due to the effects of cosmic variance on the low-ℓ multipoles, it is possible to use high-pass filters to decrease this effect and investigate how sensitive the resulting statistics are. This work has in fact already been largely done, but due to serious hardware failure it remains incomplete at the time of submission.

Part II

Exploring the local Universe with VESPA

Chapter 4

VESPA

As we have reviewed in section 1.4.6, galaxy formation and evolution are still far from being well understood. Galaxies are extremely complex objects, formed via complicated non-linear processes, and any approach (be it semi-analytical or computational) inevitably relies on simplifications. For this reason, this is an area very much driven by observations. This Chapter introduces VErsatile SPectral Analysis (VESPA) - an algorithm which recovers robust star formation and metallicity histories from galactic spectra.

4.1 Background

The stellar mass of a galaxy has been shown to correlate with properties such as luminosity, morphology, star formation rate, mass density and stellar age, to name only a few (e.g. Bell et al. 2007; Zheng et al. 2007; Brinchmann and Ellis 2000; Borch et al. 2006; Bell et al. 2003). Knowing how these relations evolve with redshift has been the goal of many observational studies, in an attempt to understand the main physical processes that drive star formation in galaxies. They can also provide strong constraints for models - these are normally "tuned" for the local Universe, and seeing how well they predict the evolution of these quantities with redshift is a very powerful test.

4.1.1 The cosmic star formation history

Measuring and understanding the star formation history of the Universe is therefore essential to our understanding of galaxy evolution - when, where and in what conditions did stars form throughout cosmic history? The traditional and simplest way to probe this is to measure the observed instantaneous star formation rate in galaxies at different redshifts. This can be achieved by looking at light emitted by young stars in the ultra-violet (UV) band or its secondary effects (e.g. Madau et al. 1996; Kennicutt 1998; Hopkins et al. 2000; Bundy et al. 2006; Erb et al. 2006; Abraham et al. 2007; Noeske et al. 2007; Verma et al. 2007). A complementary method is to look at present day galaxies and extract their star formation history, which spans the lifetime of the galaxy. Different teams have analysed a large number of galaxies in this way, whether by using the full available spectrum (Glazebrook et al. 2003b; Panter et al. 2003; Cid Fernandes et al. 2004; Heavens et al. 2004; Mathis et al. 2006; Ocvirk et al. 2006; Panter et al. 2007; Cid Fernandes et al. 2007), or by concentrating on particular spectral features or indices (e.g. Kauffmann et al. 2003; Tremonti et al. 2004; Gallazzi et al. 2005; Barber et al. 2006), which are known to be correlated with age or metallicities (e.g. Worthey 1994; Thomas et al. 2003). A compilation of the recent measurements can be seen in Figure 4.1, from Hopkins and Beacom (2006). In spite of the large scatter, it is clear from Figure 4.1 that there is a clear peak at $z \approx 2-3$, and the star formation history of the Universe has been steeply declining since. As we will see next, this sharp decline seems to be associated with high-mass galaxies.

4.1.2 Downsizing

An interesting recent result is the idea of downsizing (Cowie et al. 1996), first defined as a "smooth downward evolution in the maximum luminosity of rapidly star-forming galaxies", i.e. intense star formation is associated with increasingly less massive galaxies with decreasing redshift. There is good observational support for this trend (e.g. Cowie et al. 1996; Bell et al. 2005; Papovich et al. 2006).

However, the word's meaning has become somewhat less defined with time. Another interpretation of the word downsizing relates to the *epoch* of dominant mass build-up, which has been shown to be earlier for massive galaxies than for less massive ones (e.g. Zheng et al. 2007; Panter et al. 2007; Thomas et al. 2005; Gallazzi et al. 2005; Kauffmann et al. 2003). This effect can be clearly seen in Figure 4.2, from Panter et al.

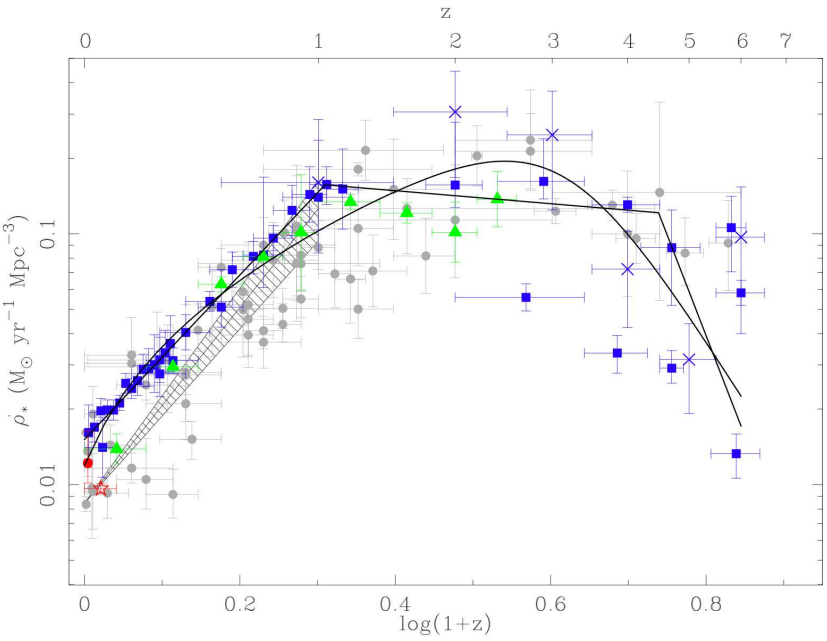

Figure 4.1: From Hopkins and Beacom (2006): the cosmic star formation history.

(2007). The lower panel of the same figure also shows that today star formation is dominated by low-mass galaxies. A similar trend was found by Sheth et al. (2006), who using clustering statistics showed that galaxies in the densest regions formed their stars earlier.

It is important to note that the former definition does not imply the latter. The former implies that star formation in massive galaxies has been decreasing with time, which culminates in a nearly complete lack of star formation in massive galaxies today. The latter simply states that mass assembly began earlier for galaxies which are massive today.

There is an apparent contradiction with the hierarchical model of galaxy formation, which suggests that if the largest galaxies are the last ones to form, then their star formation rate should be most active now. However, the largest objects, which sit in the densest peaks, are also the ones which start to assemble earlier. This reconciles our second definition of downsizing with the hierarchical model without much problem. But the first definition essentially seems to tell us that there must be a mechanism to quench star formation in massive galaxies. This must be due to what has been called

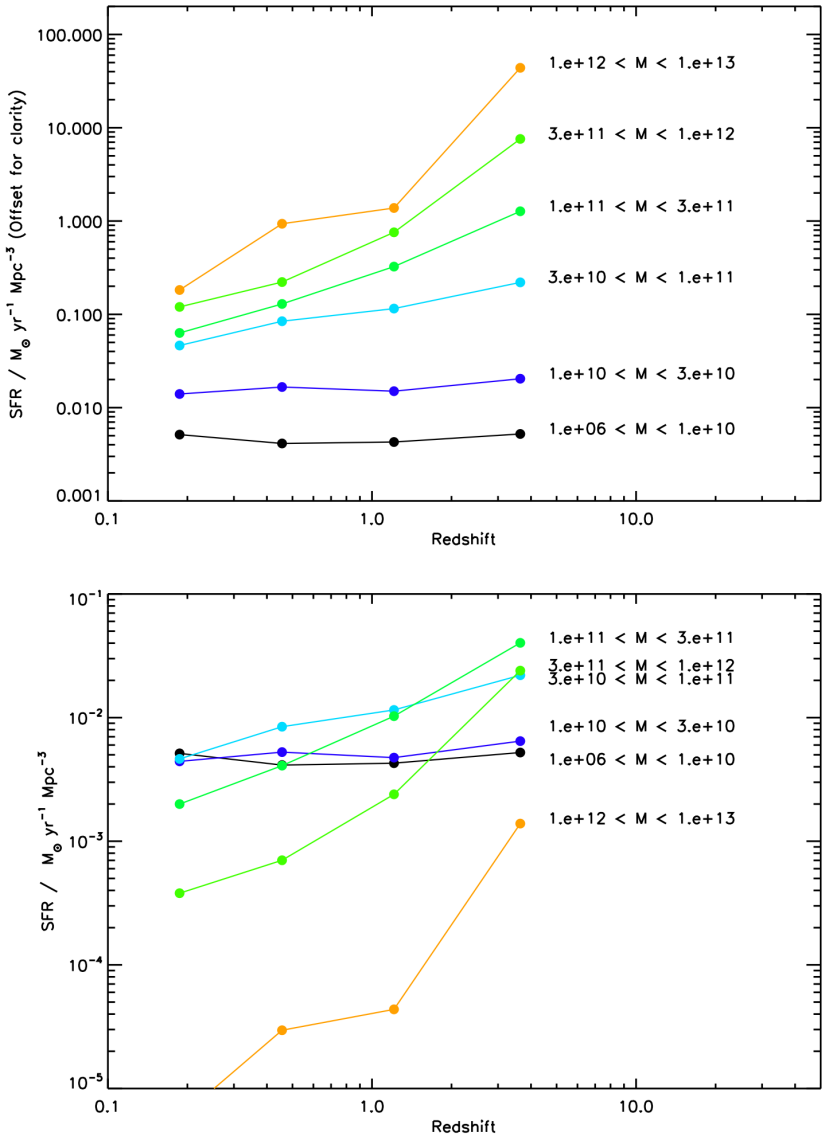

Figure 4.2: From Panter et al. (2007): the star formation rate of galaxies, estimated from the fossil record, as a function of *present-day* stellar mass. In the top panel the lines have been offset for clarity - we can see clearly that most massive stars formed their stars earlier. We can also see from the bottom panel that star formation today is dominant in low mass galaxies.

gastrophysics, a term which aims to separate the evolution of the dark matter potential wells from the less well understood gas dynamics and star formation, to which it refers.

It is easy to think of a number of mechanisms that might cause this, for example: active galactic nuclei (AGN) feedback, a depletion of gas reservoirs or a change in the infall rate of gas, a decrease in the merging rate of galaxies, or supernova feedback. Each of these would leave a different imprint in the distribution of star formation across galaxies (as a function of mass, environment, metallicity, or others - e.g. Bower et al. 2006), so again keen observation is the key here.

4.1.3 Stellar mass assembly

Integrating the cosmic star formation history of the Universe up to any given redshift (taking into account the fact that some stellar material will have been recycled, in the form of winds or star death) gives us the total stellar mass density of the Universe. Current estimates vary, but Panter et al. (2007) estimate that this number is 0.00174 in units of the critical density. An immediate thought is that this is very low compared to the estimates of Ω_b from for example WMAP5 constraints, which estimates $\Omega_b = 0.046 \pm 0.0012$. This begs the question: where have the rest of the baryons gone? Such an answer is not the immediate concern of this thesis, but one should keep in mind that there is a lot more to the baryonic Universe than stars.

A measure of how stellar mass evolves as a function of redshift provides, in principle, the same information as the cosmic star formation history. However, direct measurements of stellar mass densities at different redshifts (obtained by integrating the galaxy stellar mass function over mass) show a surprising amount of disagreement with the numbers obtained from integrating the star formation history over time. Figure 4.3 from Wilkins et al. (2008) shows this disagreement quite clearly.

The two measurements, although formally identical, are in practice obtained by very different methods. Star formation rates are more sensitive to young stars, whereas mass estimates are biased towards old stars. Different methods are also sensitive to different systematic errors and it is possible that, for example, our understanding of dust corrections for young stars at higher redshifts is simply not good enough. This discrepancy gains particular interest within the context of the IMF. The two different

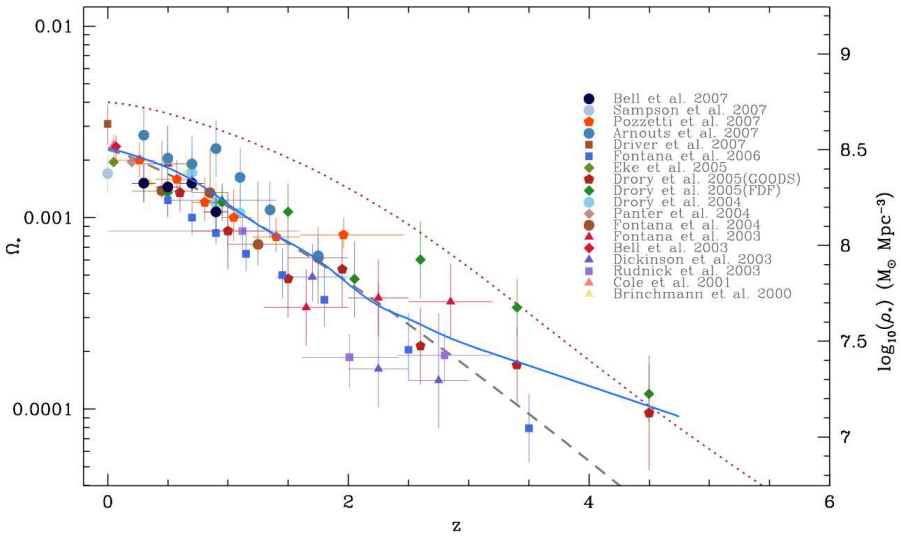

Figure 4.3: From Wilkins et al. (2008): a compilation of stellar mass density measurements. The dotted line that lies above the data points is the resulting stellar mass density from integrating the star formation history of Hopkins and Beacom (2006). There is clear disagreement between the two lines.

measurements are sensitive to mass ranges, and there would be an IMF which would bring the two measurements into agreement.

A step up from the stellar mass density, in terms of information, is of course the galaxy stellar mass function. The evolution of the shape of the galaxy stellar mass function is dictated by two main processes: the halo mass function and star formation processes. The former evolves through gravitational collapse of the dark matter and mergers, and the second through gas dynamics. A galaxy formation model must be able to explain these two processes and their change with redshift.

Some of these ideas are explored further in Chapter 5.

4.1.4 The fossil record

If we try to analyse a galaxy's luminous output in terms of a history parametrized by some chosen physical quantities, a simplification is in order. The reason is two-fold: firstly we are limited by our knowledge and ability to model all the physical processes which happen in a galaxy and produce the observed spectrum we are analysing; sec-

ondly, the observed spectrum is inevitably perturbed by noise, which intrinsically limits the amount of information we can recover.

To do this, we rely on synthetic stellar population models to describe a galaxy in terms of its stellar components, but by modelling a galaxy in this way we are intrinsically limited by the quality of the models. There are also potential concerns with flux calibration errors. However, using the full spectrum to recover the fossil record of a galaxy - or of an ensemble of galaxies - is an extremely powerful method, as the quality and amount of data relating to local galaxies vastly outshines that which concerns high-redshift galaxies. Splitting a galaxy into simple stellar populations of different ages and metallicities is a natural way of parameterising a galaxy, and it allows realistic fits to real galaxies (e.g. Bruzual and Charlot 2003). Galactic archeology has become increasingly popular in the literature recently, largely due to the increase in sophistication of stellar population synthesis codes and the improvement of the stellar spectra libraries upon which they are based, and also due to the availability of large well-calibrated spectroscopic databases, such as the Sloan Digital Sky Survey (SDSS) (York et al. 2000b; Strauss et al. 2002a).

In any case, without imposing any constraints on the allowed form of the star formation history, or perhaps an age-metallicity relation, the parameter space can become unsustainably large for a traditional approach. Ideally, one would like to do without such pre-constraints. Recently, different research teams have come up with widely different solutions for this problem. MOPED (Heavens et al. 2000) and STARLIGHT (Cid Fernandes et al. 2004) explore a well-chosen parameter space in order to find the best possible fit to the data. In the case of MOPED, this relies on compression of the full spectrum to a much smaller set of numbers which retains all the information about the parameters it tries to recover; STARLIGHT on the other hand, searches for its best fit using the full spectrum with a Metropolis algorithm. STECMAP (Ocvirk et al. 2006) solves the problem using an algebraic least-squares solution and a well-chosen regularization to keep the inversions stable. All of these and other methods acknowledge the same limitation - noise in the data and in the models introduces degeneracies into the problem which can lead to unphysical results. MOPED, for example, has produced some remarkable results concerning the average star formation history of the Universe by analysing a large sample of galaxies. However, MOPED's authors have cautioned against over-interpreting the results on a galaxy-by-galaxy basis, due to the problem

mentioned above. This is directly related to the question of how finely one should parameterise a galaxy, and what the consequences of this might be.

4.1.5 VESPA

Much of the motivation for VESPA came from the realisation that this problem will vary from galaxy to galaxy, and that the method of choosing a single parametrization to analyse a large number of galaxies can be improved on.

VESPA is based on three main ideas, which we present here and develop further in the rest of this Chapter:

- There is only so much information one can safely recover from any given set of data, and the amount of information which can be recovered from an individual galaxy varies.
- The recovered star formation fractions should be positive.
- Even though the full unconstrained problem is non-linear, it is piecewise linear in well-chosen regions of parameter space.

VESPA's ultimate goal is to derive robust information for each galaxy individually, by adapting the number of parameters it recovers on a galaxy-by-galaxy basis and increasing the resolution in parameter space only where the data warrant it. In a nutshell, this is how VESPA works: we estimate how many parameters we can recover from a given spectrum, given its noise, shape, spectral resolution and wavelength range using an analysis given by Ocvirk et al. (2006). In that paper, Singular Value Decomposition (SVD) is used to find a least squares solution, and this solution is analysed in terms of its singular vectors. VESPA uses this method only as an analysis of the solution, and uses Bounded-Variable Least-Squares (BVLS) (Stark and Parker 1995) to reach a non-negative solution in several regimes where linearity applies.

This chapter is organised as follows: in Section 4.2 we present the method, in Section 4.3 we apply VESPA to a variety of synthetic spectra, in Section 4.4 we apply VESPA to a sample of galaxies from the Sloan Digital Sky Survey (SDSS) spectroscopic database and we compare our results to those obtained with MOPED, and finally in Section 4.5

we present some final remarks on VESPA. Scientific results are presented in Chapter 5 and 6.

4.2 Method

In this section we lay down the problem to solve in detail, and explain the different steps VESPA uses to find a solution for each galaxy.

4.2.1 The problem

We assume a galaxy is composed of a series of simple stellar populations (SSPs) of varying ages and metallicities. The unobscured rest frame luminosity per unit wavelength of a galaxy can then be written as

$$F_\lambda = \int_0^t \psi(t) S_\lambda(t, Z) dt \tag{4.1}$$

where $\psi(t)$ is the star formation rate (solar masses formed per unit of time) and $S_\lambda(t, Z)$ is the luminosity per unit wavelength of a single stellar population of age t and metallicity Z, per unit mass. The dependency of the metallicity on age is unconstrained, turning this into a non-linear problem.

In order to solve this problem, we start by discretizing in wavelength and time, by averaging these two quantities into well chosen bins. For now we present the problem with a generalised parametrization, and discuss our choice in Section 4.2.3. We will use greek indices to indicate time bins, and roman indices to indicate wavelength bins.

The problem becomes

$$F_j = \sum_\alpha x_\alpha G(Z_\alpha)_{\alpha j} \tag{4.2}$$

where $F_j = (F_1, ..., F_D)$ is the luminosity of the jth wavelength bin of width $\Delta\lambda$, $G(Z_\alpha)_{\alpha j}$ is the jth luminosity point of a stellar population of age $t_\alpha = (t_1, ..., t_S)$ (spanning an age range of Δt_α) and metallicity Z_α, and $x_\alpha = (x_1, ..., x_S)$ is the total mass of population $G(Z)_{\alpha j}$ in the time bin Δt_α.

Although the full metallicity problem is non-linear, interpolating between tabulated values of Z gives a piecewise linear behaviour:

$$G(Z_\alpha)_{\alpha j} = g_\alpha G(Z_{a,\alpha})_{\alpha j} + (1 - g_\alpha) G(Z_{b,\alpha})_{\alpha j}, \tag{4.3}$$

and the problem then becomes

$$F_j = \sum_\alpha x_\alpha \left[g_\alpha G(Z_{a,\alpha})_{\alpha j} + (1 - g_\alpha)\, G(Z_{b,\alpha})_{\alpha j} \right] \tag{4.4}$$

where $G(Z_{a,\alpha})_{\alpha j}$ and $G(Z_{b,\alpha})_{\alpha j}$ are equivalent to $G(Z_\alpha)_{\alpha j}$ as above, but at fixed metallicities Z_a and Z_b, which bound the true Z. If this interpolation matches the models' resolution in Z, then we are not degrading the models in any way.

Solving the problem then requires finding the correct metallicity range. One should not underestimate the complexity this implies - trying all possible combination of consecutive values of Z_a and Z_b in a grid of 16 age bins would lead to a total number of calculations of the order of 10^9, which is unfeasible even in today's fast personal workstations. We work around this problem using an iterative approach, which we describe in Section 4.2.3.

Dust extinction

An important component when describing the luminous output of a galaxy is dust, as different wavelengths are affected in different ways. The simplest possible approach is to use one-parameter dust model, according to which we apply a single dust screen to the combined luminosity of all the galactic components. Equation (4.1) becomes

$$F_\lambda = f_{dust}(\tau_\lambda) \int_0^t \psi(t) S_\lambda(t, Z) dt \tag{4.5}$$

where we are assuming the dust extinction is the same for all stars, and characterised by the optical depth, τ_λ.

However, it is also well known that very young stars are likely to be more affected by dust. In an attempt to include this in our modelling, we follow the two-parameter dust model of Charlot and Fall (2000) in which young stars are embebbed in their birth cloud up to a time t_{BC}, when they break free into the inter-stellar medium (ISM):

$$F_\lambda = \int_0^t f_{dust}(\tau_\lambda, t) \psi(t) S_\lambda(t, Z) dt \tag{4.6}$$

and

$$f_{dust}(\tau_\lambda, t) = \begin{cases} f_{dust}(\tau_\lambda^{ISM}) f_{dust}(\tau_\lambda^{BC}), t \leq t_{BC} \\ f_{dust}(\tau_\lambda^{ISM}), t > t_{BC} \end{cases} \tag{4.7}$$

where τ_λ^{ISM} is the optical depth of the ISM and τ_λ^{BC} is the optical depth of the birth cloud. Following Charlot and Fall (2000), we take $t_{BC} = 0.03$ Gyrs.

There is a variety of choices for the form of $f_{dust}(\tau_\lambda)$. To model the dust in the ISM, we use the mixed slab model of Charlot and Fall (2000) for low optical depths ($\tau_V \leq 1$), for which

$$f_{dust}(\tau_\lambda) = \frac{1}{2\tau_\lambda}[1 + (\tau_\lambda - 1)\exp(-\tau_\lambda) - \tau_\lambda^2 E_1(\tau_\lambda)] \tag{4.8}$$

where E_1 is the exponential integral and τ_λ is the optical depth of the slab. This model is known to be less accurate for high dust values, and for optical depths greater than one we take a uniform screening model with

$$f_{dust}(\tau_\lambda) = \exp(-\tau_\lambda). \tag{4.9}$$

We only use the uniform screening model to model the dust in the birth cloud and we use $\tau_\lambda = \tau_V(\lambda/5500\mathring{A})^{-0.7}$ as our extinction curve for both environments.

As described, dust is a non-linear problem. In practice, we solve the linear problem described by equation (4.4) with a number of dust extinctions applied to the matrices $G(Z)_{ij}$ and choose the values of τ_V^{ISM} and τ_V^{BC} which result in the best fit to the data.

We initially use a binary chop search for $\tau_V^{ISM} \in [0, 4]$ and keep τ_V^{BC} fixed and equal to zero, which results in trying out typically around nine values of τ_V^{ISM}. If this initial solution reveals star formation at a time less than t_{BC} we repeat our search on a two-dimensional grid, and fit for τ_V^{ISM} and τ_V^{BC} simultaneously. There is no penalty except in CPU time to apply the two-parameter search, but we find that this procedure is robust (see section 3.4).

4.2.2 The solution

In this section we describe the method used to reach a solution for a galaxy, given a set of models and a generalised parametrization. The construction of these models and choice of parameters is described in Sections 4.2.3 and 4.2.4.

We re-write the problem described by equation (4.4) in a simpler way

$$F_j = \sum_{\kappa=1}^{2S} c_\kappa A_{\kappa j}(Z_\kappa) \tag{4.10}$$

where $Z_\kappa = Z_a$ for $\kappa < S$ and $Z_\kappa = Z_b$ for $\kappa \geq S$. $\mathbf{A}$ is a $D \times 2S$ matrix composed of synthetic models at the corresponding metallicities, and $\mathbf{c} = (c_1, ..., c_{2S})$ is the solution vector, from which the x_α and g_α in equation (4.4) can be calculated. We can then calculate a linearly interpolated metallicity at age t_α

$$Z_\alpha = g_\alpha Z_a + (1 - g_\alpha) Z_b. \tag{4.11}$$

For every age t_α we aim to recover two parameters: x_α - the total mass formed at that age (within a period Δt_α) - and Z_α - a mass-weighted metallicity.

At this stage we are not concerned with our choice of t_α and Δt_α - although these are crucial and will be discussed later. For a given set of chosen parameters, we find $\mathbf{c}$, such that

$$\chi^2 = \frac{(F_j - \sum_\kappa c_\kappa A_{\kappa j})^2}{\sigma_j^2} \tag{4.12}$$

is minimised (where σ_j is the error in the measured flux bin F_j).

A linear problem with a least squares constraint has a simple analytic solution which, for constant $\sigma_\mathbf{j}$ (white-noise) is

$$\mathbf{c_{LS}} = (\mathbf{A^T} \cdot \mathbf{A})^{-1} \cdot \mathbf{A^T} \cdot \mathbf{F} \tag{4.13}$$

In principle, any matrix inversion method, e.g. Singular Value Decomposition (SVD), can be used to solve (4.13). However, we would like to impose positivity constraints on the recovered solutions. Negative solutions are unphysical, but unfortunately common in a problem perturbed by noise.

BVLS and positivity

We use Bounded-Variable Least-Squares (BVLS) (Stark and Parker 1995) in order to solve equation (4.13). BVLS is an algorithm which solves linear problems whose variables are subject to upper and lower constraints. It uses an active set strategy to choose sets of free variables, and uses QR matrix decomposition to solve the unconstrained least-squares problem of each candidate set of free variables using (4.13):

$$\mathbf{c_{LS}} = (\mathbf{E^T} \cdot \mathbf{E})^{-1} \cdot \mathbf{E^T} \cdot \mathbf{F} \tag{4.14}$$

where $\mathbf{E}$ is effectively composed of those columns of $\mathbf{A}$ for which c_k is unconstrained, and of zero vectors for those columns for which c_k is set to zero. BVLS is an extension

of the Non-Negative Least Squares algorithm (Lawson and Hanson 1974), and they are both proven to converge in a finite number of iterations. Positivity is the only constraint in VESPA's solution.

BVLS and positivity have various advantages. Most obvious is the fact that we do away with negative solutions. In a non-constrained method (such as SVD) negative values are a response to the fact that the data is noisy. Similarly, we find that zero values returned by BVLS (in, for example, a synthetic galaxy with continuous star formation across all time) are also an artifact from noisy data. It should be kept in mind that, if the method is unbiased, this problem is solved by analysing a number of noisy realisations of the original problem - what we find is that the true values of the parameters we try to recover are consistent with the distributions yielded by this process. In this sense, not even a negative value presents a problem necessarily, as long as it is consistent with zero (or the correct solution). Although formally one might expect a solution bias from imposing positivity, empirically we have found it to be negligible when using BVLS. We therefore feel it is an advantage to discard a priori solutions we know to be unphysical.

Another advantage to using BVLS is the fact that, by fixing some parameters to the lower boundary (zero, in this case), it effectively reduces the number of fitting parameters to the number of those which keeps unconstrained. Given the overall aims of VESPA, this has proven to be advantageous.

Noise

The inversion in equation (4.13) is often highly sensitive to noise, and care is needed when recovering solutions with matrix inversion methods. The fit in data-space will always improve as we increase the number of parameters, but these might not all provide meaningful information. We follow an analysis given in Ocvirk et al. (2006) in order to understand how much this affects our results, and to choose a suitable age parametrization for each galaxy. This is not an exact method, and it does not guarantee that the solutions we recover have no contribution from noise. However, we found that in most cases it provides a very useful guideline (see section 4.3.3, in particular Figure 4.15).

We refer the reader to the above paper for a full discussion, and we reproduce here the steps used in our analysis.

We use SVD to decompose the model matrix $\mathbf{E}$ as

$$\mathbf{E} = \mathbf{U} \cdot \mathbf{W} \cdot \mathbf{V^T} \tag{4.15}$$

where $\mathbf{U}$ is a $D \times 2S$ orthonormal matrix with singular data-vectors $\mathbf{u}_\kappa$ as columns, $\mathbf{V}$ is a $2S \times 2S$ orthonormal matrix with the singular solution-vectors $\mathbf{v}_\kappa$ as columns, and $\mathbf{W}$ is a $2S \times 2S$ diagonal matrix $\mathbf{W} = diag(w_1, ..., w_{2S})$ where w_κ are the matrix singular values in decreasing order. Replacing $\mathbf{E}$ by this decomposition in equation (4.13) gives

$$\mathbf{c}_{LS} = \mathbf{V} \cdot \mathbf{W}^{-1} \cdot \mathbf{U}^T \cdot \mathbf{F} = \sum_{\kappa=1}^{2S} \frac{\mathbf{u}_\kappa^T \cdot \mathbf{F}}{w_\kappa} \mathbf{v}_\kappa \tag{4.16}$$

The solution vector is a linear combination of the solution singular values, parametrized by the dot product between the data and the corresponding data singular vector, and divided by the k^{th} singular value. The data vector itself is a combination of the true underlying emitted flux and noise: $\mathbf{F} = \mathbf{F_{true}} + \mathbf{e}$. Equation (4.16) becomes

$$\mathbf{c}_{LS} = \sum_{\kappa=1}^{2S} \frac{\mathbf{u}_\kappa^T \cdot \mathbf{F}_{true}}{w_\kappa} \mathbf{v}_\kappa + \sum_{\kappa=1}^{2S} \frac{\mathbf{u}_\kappa^T \cdot \mathbf{e}}{w_\kappa} \mathbf{v}_\kappa \equiv \mathbf{c}_{true} + \mathbf{c}_e \tag{4.17}$$

where $\mathbf{c_{true}}$ is the solution vector to the noiseless problem and $\mathbf{c_e}$ is an unavoidable added term due to the presence of noise.

It is extremely informative to compare the amplitudes of the two terms in the sum (4.17), and to monitor their contributions to the solution vector with varying rank. In Figure 4.4 we plot $|\mathbf{u_\kappa^T} \cdot \mathbf{F}|$ and $\mathbf{u_\kappa^T} \cdot \mathbf{e}$ as a function of rank κ, for a synthetic spectrum with a SNR per pixel of 50 (at a resolution of 3Å) and an exponentially-decaying star formation history. We observe the behaviour described and discussed in Ocvirk et al. (2006). The combinations associated with the noise terms maintain a roughly constant power across all ranks, with a an average value of $\langle \mathbf{F} \rangle$ /SNR. The data terms, however, drop significantly with rank, and we can therefore identify two ranges: a noise-dominated κ-range, in which the noise contributions match or dominate the true data contributions, and a data-dominated range, where the contributions to the solution are largely data motivated. We call the transition rank κ_{crit}. Overall, high-κ ranks tend to dominate the solution, since the singular values w_κ decrease with κ. This only amplifies the problem by giving greater weight to noise-dominated terms in the sum (4.16). Figure 4.5 shows the contribution coming from each rank κ to the final solution

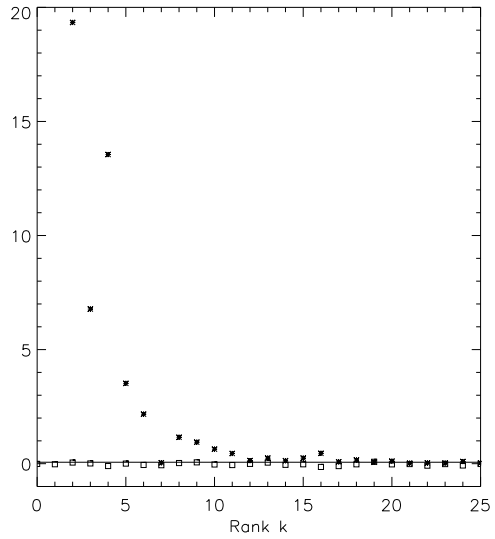

Figure 4.4: The behaviour of the singular values with matrix rank k. The stars are $|\mathbf{u}_\kappa^T \cdot \mathbf{F}|$ and the squares are $\mathbf{u}_\kappa^T \cdot \mathbf{e}$. The line is $\langle \mathbf{F} \rangle$ /SNR, which in this case has a value of approximately 0.06.

- the coefficient $(\mathbf{u}_\kappa^T \cdot \mathbf{F})/w_\kappa$. We see this weight increases with rank.

Whereas this analysis gives us great insight into the problem, we do not in fact use the sum (4.16) to obtain $\mathbf{c}_{LS}$, for the reasons given in section 4.2.2.

For real data we are only able to calculate $\mathbf{u}_\kappa^T \cdot \mathbf{F}$ and estimate the noise level at $\langle \mathbf{F} \rangle$ /SNR and we use this information to estimate the number of non-zero parameters to recover from the data. Our aim is to a have a solution which is dominated by the signal, and not by the noise. We therefore want our number of non-zero recovered parameters to be less than or equal to κ_{crit}. Estimating where this transition happens is always a noisy process. In this thesis we take the conservative approach of setting κ_{crit} to be the rank at which the perturbed singular values first cross the $\langle \mathbf{F} \rangle$ /SNR barrier. In the case of Figure 4.4 this happens at $\kappa_{crit} = 7$.

4.2.3 Choosing a galaxy parametrization

One of the advantages of VESPA is that it has the ability to choose the number of parameters to recover in any given galaxy. This is possible due to a time grid of varying

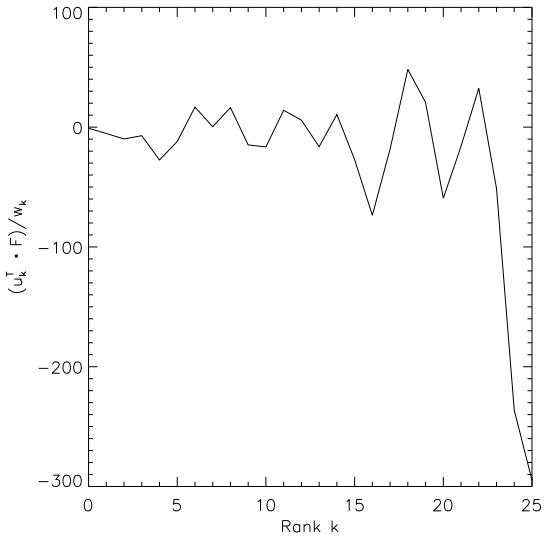

Figure 4.5: The coefficients in sum (4.16) as a function of rank κ. We see that the highest rank modes (corresponding to the smaller singular values) tend to contribute the most to the solution.

resolutions, which VESPA can explore to find a solution. This section describes this grid and the criteria used to reach a final parametrization.

The grid

We work on a grid with a maximum resolution of 16 age bins, logarithmically spaced in lookback time from 0.02 up to 14 Gyr. The grid has three further resolution levels, where we split the age of the Universe in eight, four and finally two age bins, also logarithmically spaced in the same range. A schematic view of this grid can be seen in Figure 4.6.

The idea behind the multi-resolution grid is to start our search with a low number of parameters (in coarser resolution, so that the entirety of the age of Universe is covered), and then increase the resolution only where the data warrant it by splitting the bin with the highest flux contribution in two, and so on. In effect, we construct one such grid for each of the tabulated metallicities, Z_a and Z_b. We work with five metallicity values, $Z = [0.0004, 0.004, 0.008, 0.02, 0.05]$ which correspond to the metallicity resolution of

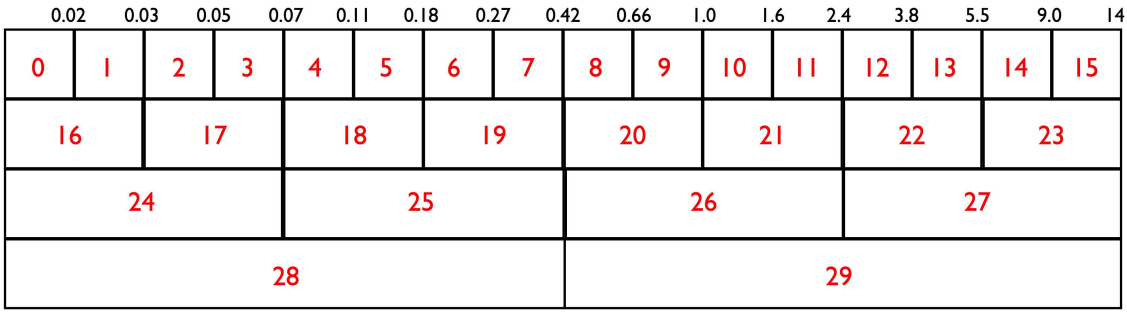

Figure 4.6: Schematic view of the grid of bins used by VESPA. The top line of black numbers indicates the age of each boundary. The red numbers in each of the bins is an unique bin identifier number, which can be used to quickly retrieve properties of a given bin (see Chapter 5 for more information).

the models used, where Z is the fraction of the mass of the star composed of metals ($Z_{\odot} = 0.02$). The construction of the models for each of the time bins is discussed in Section 4.2.4.

To each of the grids we can apply a dust extinction as explained in Section 4.2.1.

The search

We go through the following steps in order to reach a solution:

1. We begin our search with three bins: two bins of width 4 and one bin of width 8 (oldest), where here we are measuring widths in units of high-resolution bins.

2. We calculate a solution using equation (4.10) for every possible combination of consecutive boundaries Z_a and Z_b, and we choose the one which gives the best value of reduced χ^2.

3. We calculate the number of perturbed singular values above the noise level, as described at the end of Section 4.2.2.

4. We find the bin which contributes the most to the total flux and we split it into two.

5. We find a solution in the new parametrization, this time by trying out all possible combinations of Z_a and Z_b for the newly split bins only, and fixing the metallicity boundaries of the remainder bins to the boundaries obtained in the previous solution. If a bin had no stars in the previous iteration, we set $Z_a = 0.0004$ and $Z_b = 0.05$.

6. We return to step 3 and we proceed until we have reached the maximum resolution in all populated bins.

7. We look backwards in our sequence of solutions for the last instance with a number of non-zero recovered parameters equal to or less than κ_{crit} as calculated in step 3 and take this as our best solution.

We illustrate this sequence in Figure 4.7, where we show the evolution of the search in a synthetic galaxy composed of two stellar bursts of equal star formation rates - one young and one old. VESPA first splits the components which contribute the most to the total flux. In this case this is the young burst which can be seen in the first bin. Even though VESPA always resolves bins with any mass to the possible highest resolution, it then searches for the latest solution which has passed the SVD criterion explained in Section 4.2.2. In this case, this corresponds to the fifth from the top solution. VESPA chooses this solution in favour of the following ones due to the number of perturbed singular values above the solid line (right panel). In this case, the solution chosen by VESPA is a better fit in parameters space (note the logarithmic scale in the y-axis - the following solution put the vast majority of the mass in the wrong bin). We observed this type of improvement in the majority of all cases studied (see Figure 4.15).

The final solution

Our final solution comes in a parametrization such that the total number of non-zero recovered parameters is less than or equal to the number of perturbed singular values

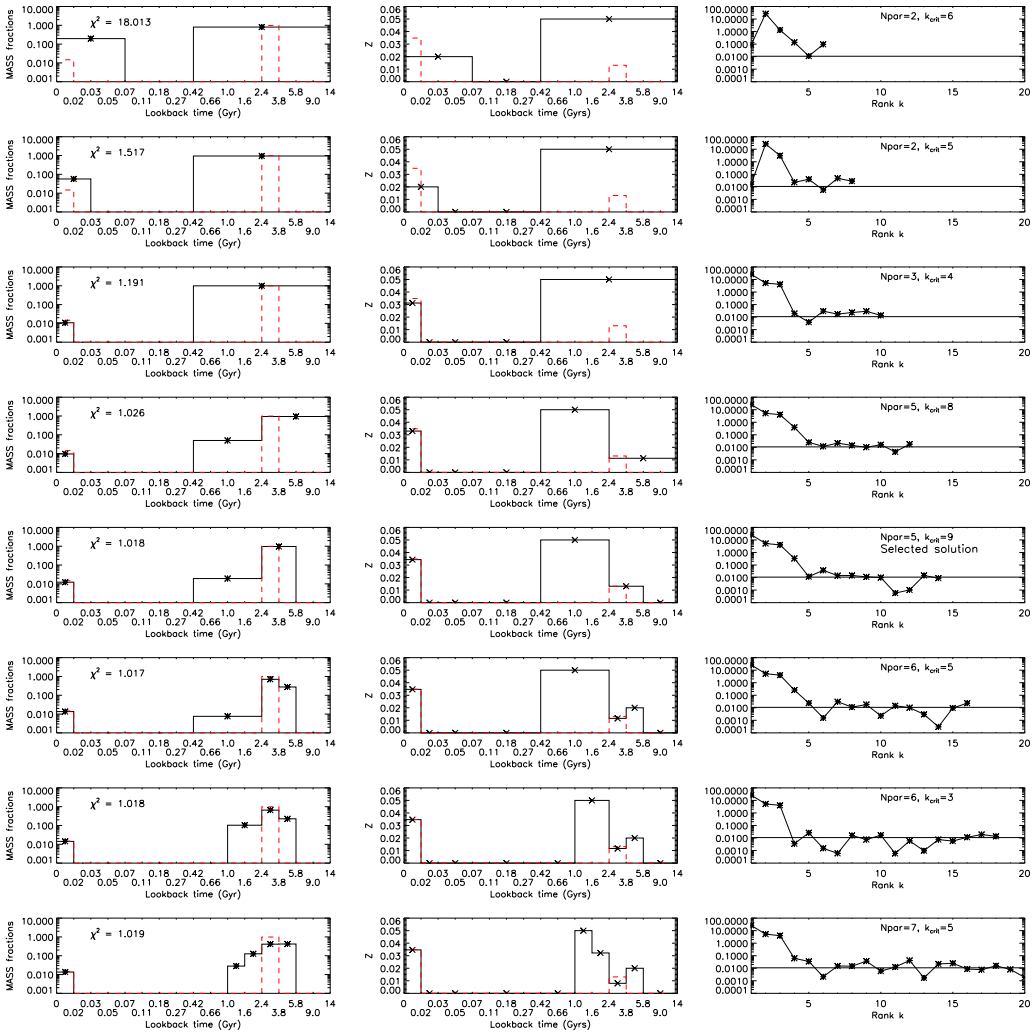

Figure 4.7: The evolution of the fit, as VESPA searches for a solution. Sequence should be read from top to bottom. Each line shows a stage in the sequence: the left panel shows the input star formation history in the dashed line, and the recovered mass fractions on the solid line for a given parametrization ; the middle panel shows the input metallicities in the dashed line, and the recovered metallicities on the solid line; the right panel shows the absolute value of the perturbed singular values $|\mathbf{u}_\kappa \cdot \mathbf{F}|$ (stars and solid line) and the estimated noise level $\langle \mathbf{F} \rangle$ /SNR. In this panel we also show the value of κ_{crit} and the number of non-zero elements of $\mathbf{c_{LS}}$ in each iteration. The chosen solution is the **fifth from the top**, and indicated accordingly. This galaxy consists of two burst events of equal star formation rate - a very young and an old burst. It was modelled with a resolution of 3Å and a signal-to-noise ratio per pixel of 50. We see the recovery is good but not perfect - there is a 1 per cent leakage from the older population - but better than the following solutions, where this bin is split. See text in Section 4.2.3 for more details.

above the estimated noise level.

The above sequence is performed for each of several combinations of $\tau_V^{BC}, \tau_V^{ISM}$, and we choose the attenuation which provides the best fit.

For each galaxy we recover N star formation masses, with an associated metallicity, where N is the total number of bins, and a maximum of two dust parameters.

4.2.4 The models

The backbone to our grid of models is the BC03 set of synthetic SSP models (Bruzual and Charlot 2003), with a Chabrier initial mass function (Chabrier 2003) and Padova 1994 evolutionary tracks (Alongi et al. 1993; Bressan et al. 1993; Fagotto et al. 1994a,b; Girardi et al. 1996). Although any set of stellar population models can be used, these provide a detailed spectral evolution of stellar populations over a suitable range of wavelength, ages and metallicities: $S(\lambda, t, Z)$. The models have been normalised to one solar mass at the age $t = 0$.

High-resolution age bins

At our highest resolution we work with 16 age bins, equally spaced in a logarithmic time scale between now and the age of the Universe. In each bin, we assume a constant star formation rate

$$f_\alpha^{HR}(\lambda, Z) = \psi \int_{\Delta t_\alpha} S(\lambda, t, Z) dt \tag{4.18}$$

with

$$\psi = 1/\Delta t_\alpha. \tag{4.19}$$

Low-resolution age bins

As described in Section 4.2.3, we work on a grid of different resolution time bins and we construct the low resolution bins using the high resolution bins described in Section 4.2.4. We do not assume a constant star formation rate in this case, as in wider bins the light from the younger components would largely dominate over the contribution from the older ones. Instead, we use a decaying star formation history, such that the

light contributions from all the components are comparable. Recall equation (4.1)

$$f_\alpha^{LR}(\lambda, Z) = \int_{\Delta t_\alpha} \psi(t) S(\lambda, t, Z) dt, \tag{4.20}$$

which we approximate to

$$f_\beta^{LR}(\lambda, Z) = \frac{\sum_{\alpha\in\beta} f_\alpha^{HR}(\lambda, Z)\psi_\alpha \Delta t_\alpha}{\sum_{\beta\in\alpha} \psi_\alpha \Delta t_\alpha} \tag{4.21}$$

where low resolution bin β incorporates the high resolution bins $\alpha \in \beta$, and we set

$$\psi_\alpha \Delta t_\alpha = \frac{1}{\int_\lambda f_\alpha^{HR}(\lambda, Z) d\lambda}. \tag{4.22}$$

Depending on the galaxy, the final solution obtained with the sequence detailed in Section 4.2.3 can be described in terms of low-resolution age bins. In this case we should interpret the recovered mass as the total mass formed during the period implied by the width of the bin, but we cannot make any conclusions as to when in the bin the mass was formed. Similarly, the recovered metallicity for the bin should be interpreted as a mass-weighted metallicity for the total mass formed in the bin.

4.2.5 Errors

The quality of our fits and of our solutions is affected by the noise in the data, the noise in the models, and the parametrization we choose (which does not reflect the complete physical scenario within a galaxy). We aim to apply VESPA firstly to SDSS galaxies, which typically have a SNR ≈ 20 per resolution element of 3Å, which puts us in a regime where the main limitations come from the noise in the data.

To estimate how much noise affects our recovered solutions we take a rather empirical approach. For each recovered solution we create n_{error} random noisy realisations and we apply VESPA to each of these spectra. We re-bin each recovered solution in the parametrization of the solution we want to analyse and estimate the covariance matrices

$$C(x)_{\alpha\beta} = \langle (x_\alpha - \bar{x}_\alpha)(x_\beta - \bar{x}_\beta) \rangle \tag{4.23}$$

$$C(Z)_{\alpha\beta} = \left\langle (Z_\alpha - \bar{Z}_\alpha)(Z_\beta - \bar{Z}_\beta) \right\rangle. \tag{4.24}$$

All the plots in Sections 4.3 and 4.4 show error bars derived from $C_{\alpha\alpha}^{1/2}$, although it is worth keeping in mind that these are typically highly correlated.

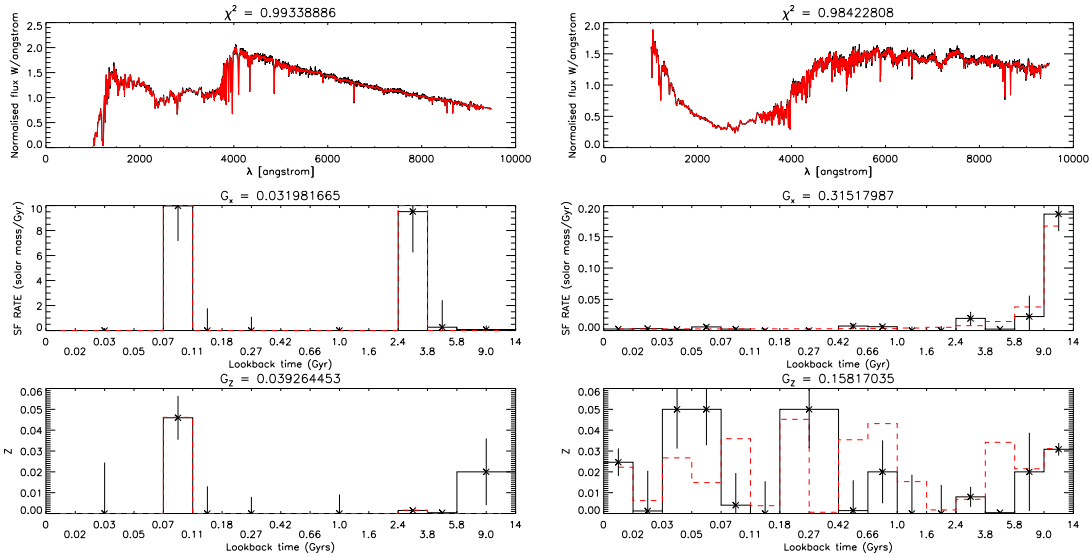

Figure 4.8: Two examples of VESPA's analysis on synthetic galaxies. The top panels show the original spectrum in the dark line and fitted spectrum in the red line. The middle panels show the input (dashed, red) and the recovered (solid, black) star formation rates and the bottom panel shows the input (dashed, red) and recovered (solid, black) metallicities per bin. Note that even though many of the recovered metallicities are wrong, these tend to correspond to bins with very little star formation, and are therefore virtually unconstrained.

4.2.6 Timings

A basic run of VESPA (which consists of roughly 5 runs down the sequence detailed in Section 4.2.3, one for each value of dust extinction) takes about 5 seconds. If accurate error estimations are needed per galaxy, this will add another one or two minutes to the timing, depending on how accurately one would like to estimate the covariance matrices, and depending on the number of data points. With $n_{error} = 10$, a typical SDSS galaxy takes around one minute to analyse.

4.3 Tests on Simulated Data

We tested VESPA on a variety of synthetic spectra, in order to understand its capabilities and limitations. In particular, we tried to understand the effect of three factors in the quality of our solutions: the input star formation history, the noise in the data, and the wavelength coverage of the spectrum. We have also looked at the effects of dust extinction. Throughout we have modelled our galaxies in a resolution of 3Å.

Even though we are aware that showing individual examples of VESPA's results from

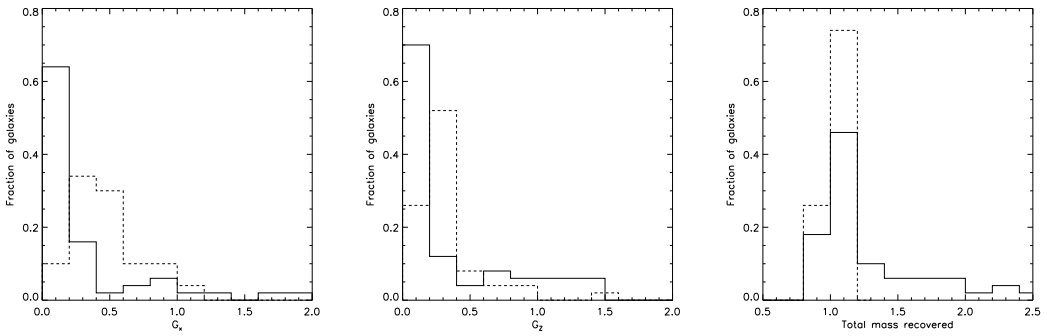

Figure 4.9: The distribution of G_x, G_Z and total mass recovered for 50 galaxies with a SNR per pixel of 50. Solid lines correspond to dual burst and dashed lines to exponentially decaying ones. See text in Section 4.3.1 for details.

synthetic spectra can be extraordinarily unrepresentative, we feel obliged to show a few for illustration purposes. We will show a typical result for most of the cases we present, but we also define some measurements of success, so that the overall performance of VESPA can be tracked as we vary any factors. We define

$$G_x = \sum_{\alpha} \left| \frac{x_\alpha - x_\alpha^I}{x_\alpha^I} \right| \omega_\alpha \tag{4.25}$$

and

$$G_Z = \sum_{\alpha} \left| \frac{Z_\alpha - Z_\alpha^I}{Z_\alpha^I} \right| \omega_\alpha \tag{4.26}$$

where x_α^I and Z_α^I are the total mass and correspondent metallicity in bin α (re-binned to match the solution's parametrization if necessary), and ω_α is the flux contribution of population of age t_α. G_x and G_Z are a flux-weighted average of the total absolute fractional errors in the solution, and give an indication of how well VESPA recovers the most significant parameters. A perfect solution gives $G_x = G_Z = 0$. It is also worth noting that this statistic does not take into account the associated error with each recovered parameter - deviations from the true solution are usually expected given the estimated covariance matrices. We will also show how these factors affect the recovered total mass for a galaxy. In all cases we have re-normalised the total masses such that total input mass for each galaxy is 1.

4.3.1 Star formation histories

We present here some results for synthetic spectra with two different star formation

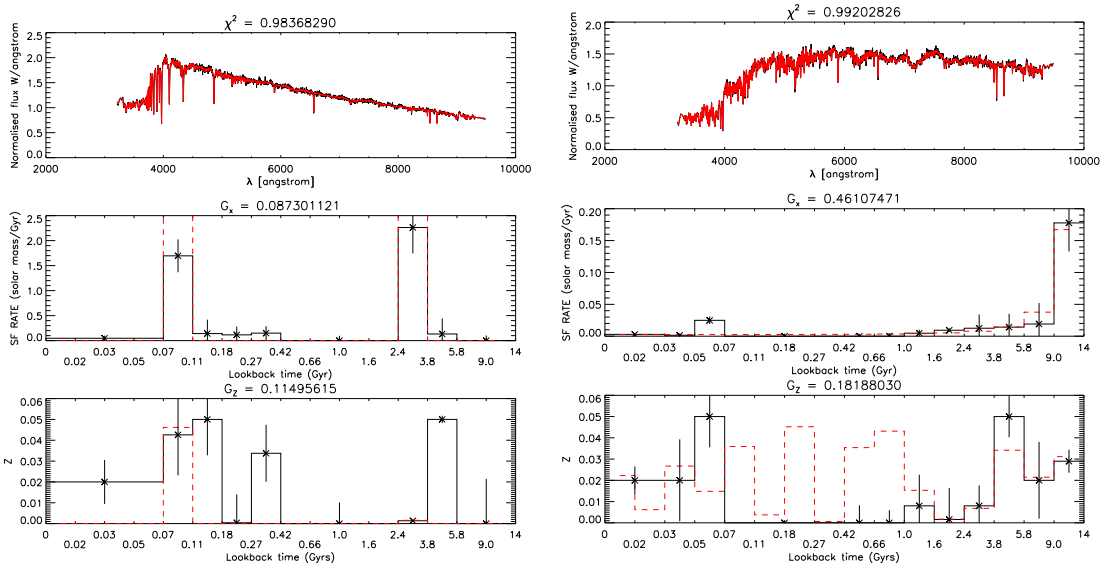

Figure 4.10: Same galaxies as in Figure 4.8, but results obtained by using a smaller wavelength range. The goodness-of-fit in data space is still excellent, but it becomes more difficult to break certain degeneracies.

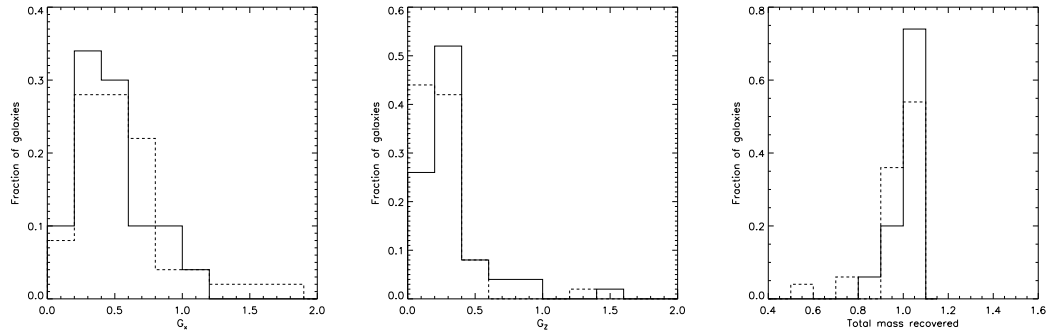

Figure 4.11: The distribution of G_x, G_Z and total mass recovered for 50 galaxies with a SNR per pixel of 50 and two different wavelength coverage. Solid line corresponds to $\lambda \in [1000, 9500]$Å and dashed line to $\lambda \in [3200, 9500]$ Å.

histories. All of the spectra in this section were synthesised with a SNR per pixel of 50, and we initially fit the very wide wavelength range $\lambda \in [1000, 9500]$Å.

We choose two very difference cases: firstly a star formation history of dual bursts, with a large random variety of burst age separations and metallicities (where we set the star formation rate to be 10 solar masses per Gyr in all bursts). Secondly, we chose a SFH with an exponentially decaying star formation rate: SFR $\propto \exp(\gamma t_\alpha)$, where t_α is the age of the bin in lookback time in Gyr. Here we show results for $\gamma = 0.3$ Gyr^{-1}. Rather than being physically motivated, our choice of γ reflects a SFH which is not too steep as to essentially mimic a single old burst, but which is also not completely dominated by recent star formation. In all cases the metallicity in each bin is randomly set. Figure 4.8 shows a typical example from each type.

Figure 4.9 shows the distribution of G_x, G_Z and of the recovered total masses for a sample of 50 galaxies. We see differences between the two cases. Firstly, in dual burst galaxies, we seem to do better in recovering data from significant individual bins, but worse in overall mass. This reflects the fact that G_x is dominated by the fractional errors in the most significant bins, but the total mass can be affected by small flux contributions in old bins which can have large masses. On the other hand, with an exponentially decaying star formation rate, we do worse overall (although this is mainly a reflection that more bins have significant contributions to the flux) but we recover the total mass of the galaxy exceptionally well.

4.3.2 Wavelength range

Wavelength range is an important factor in this sort of analysis, as different parts of the spectrum will help to break different degeneracies. Since we are primarily interested in SDSS galaxies, we have studied how well VESPA does in the more realistic wavelength range of $\lambda \in [3200, 9500]$ Å.

Figure 4.10 shows the results for the same galaxies shown in Figure 4.8, obtained with the new wavelength range. In these particular cases, we notice a more pronounced difference in the dual bursts galaxy, but looking at a more substantial sample of galaxies shows that this is not generally the case. Figure 4.11 shows G_x, G_Z and total mass recovered for 50 exponentially decaying star formation history galaxies, with a signal to noise ratio of 50 and the two different wavelength ranges. We do not see a largely

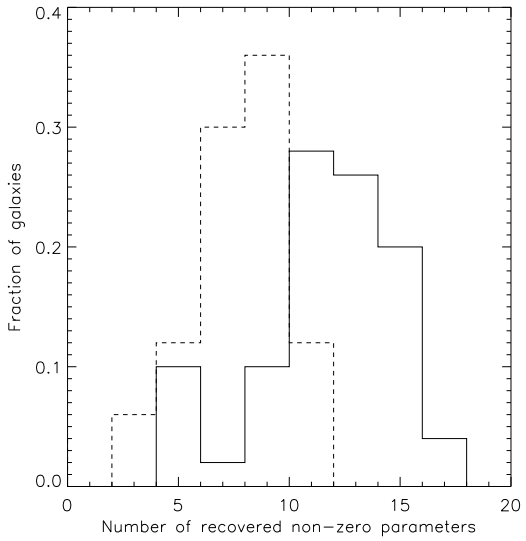

Figure 4.12: The recovered number of non-zero parameters in 50 galaxies with an exponentially decaying star formation history, using different wavelength ranges: $\lambda \in [1000, 9500]$ Å(solid line) and $\lambda \in [3200, 9500]$ Å(dashed line). Please note that these correspond to the *total* number of non-zero components in the solution vector $\mathbf{c}_\kappa$ and not to the number of recovered stellar populations.

significant change in both cases, and we observe a less significant difference in the dual bursts galaxies (not plotted here).

We find it instructive to keep track of how many parameters we recover in total, as we change any factors. Figure 4.12 shows an histogram of the total number of non-zero parameters we recovered from our sample galaxies with exponentially-decaying star formation histories and both wavelength ranges. Note that these are the components of the solution vector $\mathbf{c}_\kappa$ which are non-zero - they do not represent a number of recovered stellar populations. In this case there is a clear decrease in the number of recovered parameters, suggesting a wider wavelength range is a useful way to increase resolution in parameter space.

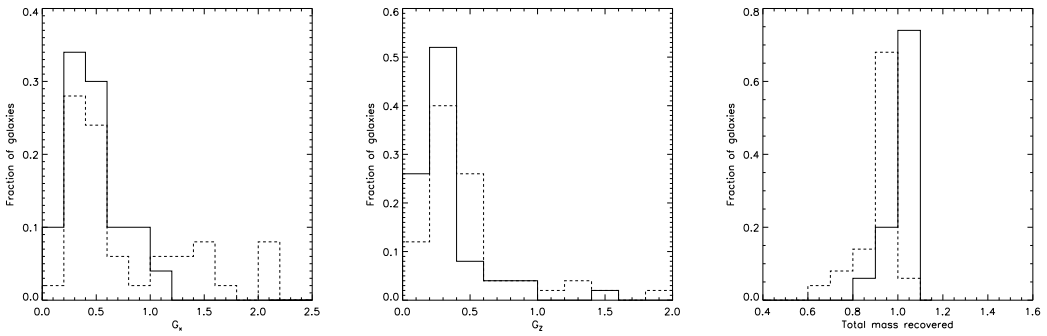

Figure 4.13: The distribution of G_x, G_Z and total mass recovered for 50 galaxies with an exponential decaying star formation history and different signal-to-noise ratios. Solid lines correspond to SNR = 50 and dashed lines to SNR =20. See text in Section 4.3.3 for details.

4.3.3 Noise

It is of interest to vary the signal-to-noise ratio in the synthetic spectra. We have repeated the studies detailed in the two previous sections with varying values of noise, and we investigate how this affects both the quality of the solutions and their resolution in parameter space.

Figure 4.14 shows how the recovered number of parameters changed by increasing the noise in the galaxies with an exponentially decaying star formation rate and wide wavelength range. In this case the increase in the noise leads to a significant reduction of the number of parameters recovered for each galaxy. This behaviour is equally clear for different star formation histories and different wavelength coverage, and is directly caused by the stopping criterion defined in Section 4.3.3.

The quality of the solutions is also affected by this increase in noise, as can be seen in figure 4.13, where we have plotted G_x, G_Z and the total recovered mass for two different values of SNR. The quality of the solutions decreases with the higher noise levels, as is to be expected. However, a more interesting question to ask is whether this decrease in the quality of the solutions would indeed be more pronounced without the SVD stopping criterion. Figure 4.15 shows a comparison between G_x obtained as we have described and obtained without any stopping mechanism (so letting our search go to the highest possible resolution and taking the final solution) for 50 galaxies with an exponentially decaying star formation history and a signal-to-noise ratio of 20. The results show clearly that there is a significant advantage in using the SVD stopping

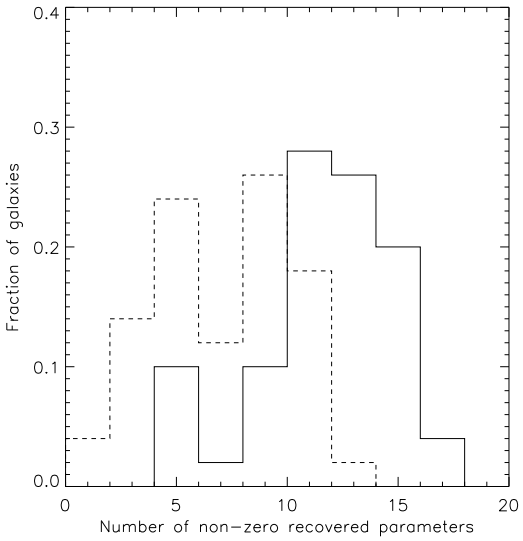

Figure 4.14: The recovered number of non-zero parameters as we change the noise in the data from 50 (solid line) to 20 (dashed line), in a sample of galaxies with an exponentially decaying star formation rate. Please note that these correspond to the *total* number of non-zero components in the solution vector $\mathbf{c}_\kappa$ and not to the number of recovered stellar populations.

criterion. Naturally, the goodness of fit in data space is consistently better as we increase the number of parameters but this improvement is illusory - the parameter recovery is worse. This is exactly the expected behaviour - we choose to sacrifice resolution in parameter space in favour of a more robust solution - even though naively one could think a lower χ^2 solution would indicate a better solution. The significance of this improvement changes with the amount of noise and wavelength range of the data (and to a lesser extent with type of star formation history) but we observed an improvement in all cases we have studied.

As expected, further decreasing the signal-to-noise ratio leads to a further degradation of the recovered solutions. This is accompanied by a suitable increase in the error bars and correlation matrices, but in cases of a SNR$\approx$ 10 and less it becomes very difficult to recover any meaningful information from individual spectra.

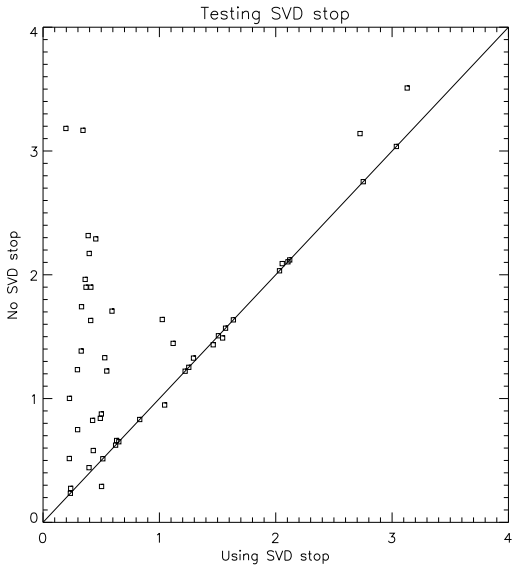

Figure 4.15: Testing the SVD stopping criterion. Plots show goodness of fit G_x for the solution of 50 galaxies obtained with and without the SVD stopping criterion. We see that by recovering only as many parameters as the data warrants gives improved parameter estimation in almost all cases, and a striking improvement in many.

4.3.4 Dust

In this section we use simulated galaxies to study the effect of dust in our solutions. As explained in Section 4.2.1, due to the non-linear nature of the problem, we cannot include dust as one of the free parameters analysed by SVD. Instead, we fit for a maximum of two dust parameters using a brute-force approach which aims to minimise χ^2 in data-space by trying out a series of values for τ_V^{ISM} and τ_V^{BC}.

For each galaxy we assign random values of $\tau_V^{ISM} \in [0, 2]$ and $\tau_V^{BC} \in [1, 2]$ and we are interested in how well we recover these parameters and any possible degeneracies.

Figure 4.16 shows the input and recovered values for τ_V^{ISM} for galaxies with a signal-to-noise ratio of 50, and which were analysed using the wavelength range $\lambda \in [3200, 9500]$Å. We show results for two different cases of star formation history: 50 galaxies with an exponentially-decaying SFR and 50 galaxies formed by dual-bursts. We observe a good recovery of τ_V^{ISM} in both cases, especially at low optical depths.

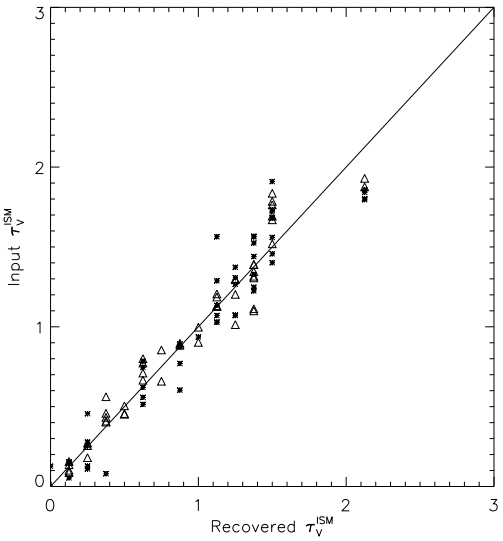

Figure 4.16: Testing the recovery of τ_V^{ISM} for 50 galaxies with a exponentially-decaying star formation history (triangles) and 50 galaxies formed with a random combination of dual bursts (stars). The input values are randomly chosen and continuously distributed between 0 and 2. The recovered values are chosen from a tabulated grid between 0 and 4.

However, we mostly observe a poor recovery of τ_V^{BC}, especially at high optical depths. This is unsurprisingly flagging up a certain level of degeneracy between mass and degree of extinction, which gets worse as the optical depth increases. Essentially, it becomes difficult to distinguish between a highly obscured massive population and a less massive population surrounded by less dust. It is worth keeping in mind that young populations are affected by both dust components simultaneously, and generally, even though the recovery of the second dust parameter may not be accurate, it allows for a better estimation of the dominant dust component.

This can be tested by simulating galaxies on a two-component dust model and by analysing them using both a single component model, and a two-component model. E.g., when using the more sophisticated model, we noted that the mean error on τ_V^{ISM} on a subsample of dual-burst galaxies (synthesised as explained in section 4.3.1, but chosen to have young star formation) was reduced from 35 to 28 per cent. This simple test also revealed that we are less likely to underestimate the mass of young populations by allowing an extra dust component, but that we are also introducing an extra

degeneracy, especially so in the case of faint young populations. However, we feel that the two-parameter dust model brings more advantages than disadvantages, with the caveat being that dusty young populations can be poorly constrained. In any case, we note that each galaxy is always analysed with a one-parameter model before being potentially analysed with a two-parameter model, and both solutions are kept and always available for analysis.

Finally, our test also partly justifies our choice to first run a single dust component model and only apply a two-component model if we detect stars in the first two bins - we find that although a one-component model might underestimate the amount of young stars, it does not fail to detect them. We repeated a similar test in real data, by analysing the same sample with and one- and a two-parameter dust model. We found similar results, with a one-parameter model failing to yield star formation in young bins only around 1 per-cent of the time (compared to the two-parameter model), and only in cases where the contribution of the light from the young populations was very small (of the order of 1 to 2 per cent).

4.4 Tests on real data

In this section we present some results obtained by applying VESPA to galaxies in the SDSS. Our aim is to analyse these galaxies, and to produce and publish a catalogue of robust star formation histories, stellar masses and dust contents, from which a wealth of information can be derived. This catalogue is described in Chapter 5. A description of the SDSS survey, data, and our pre-processing of the spectra before analysis can also be found in Chapter 5.

Here we present results from a sub-sample of galaxies, which we used to test VESPA in a variety of ways.

4.4.1 Duplicate galaxies

There are a number of galaxies in the SDSS database which have been observed more than once, for a variety of reasons. This provides an opportunity to check how variations in observation-dependent corrections affect the results obtained by VESPA.

We have used a subset of the sample of duplicate objects in Brinchmann et al. (2004)[1]

[1] Available at http://www.mpa-garching.mpg.de/SDSS/

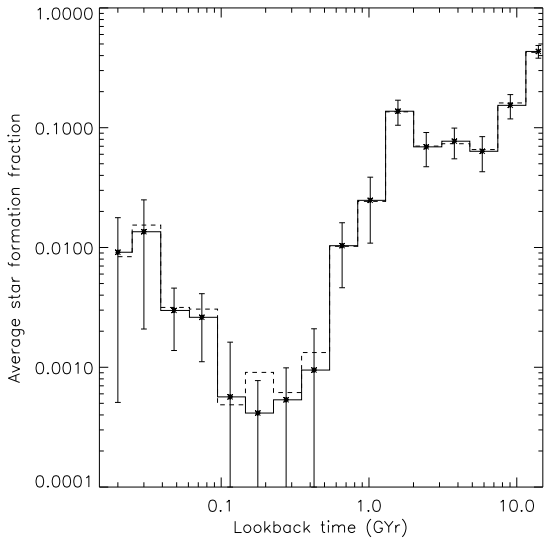

Figure 4.17: Average star formation fraction as a function of lookback time for the 2000 galaxies in list A (solid line) and list B (dashed line). The error bars shown are the errors bars on the mean for each age bin. We show only the errors from list A to avoid cluttering - the errors from list B are of similar amplitude.

to create two sets of observations for 2000 galaxies, which we named list A and list B. We are interested in seeing how the errors we estimate for our results compare to errors introduced by intrinsic variations caused by changing the observation conditions (such as quality of the spectra, placement of the fibre, sky subtraction or spectrophometric calibrations).

Figure 4.17 shows the average star formation fraction as a function of lookback time for both sets of observation. The error bars showed are errors on the mean. We see no signs of being dominated by systematics when estimating the star formation fraction of a sample of galaxies.

Figure 4.18 shows the total stellar mass obtained for a set of 500 galaxies in both observations (details of how we estimate the total stellar mass of a galaxy are included in section 4.4.3). The error bars are obtained directly from the estimated covariance matrix $C(x)$ (equation 4.23). Even though most of the galaxy duplicates produce mass estimates in agreement with each other given the error estimates, a minority does not. Upon inspection, these galaxies show significant differences in their continuum, but

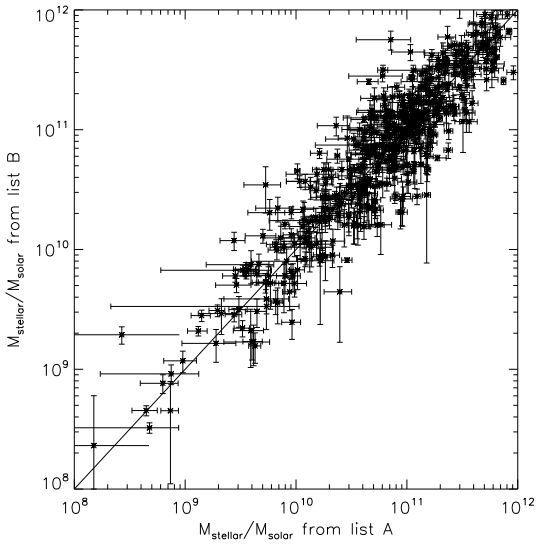

Figure 4.18: Total stellar mass recovered for two sets of observations of 500 galaxies in the main galaxy sample. The error bars are calculated from $C(x)$.

after further investigation it remains unclear what motivates such a difference. The simplest explanation is that the spectrophotometric calibration differs significantly between both observations, and that might have been the reason the plate or object was re-observed. Whatever the reason however, the clear conclusion is that stellar mass estimates are highly sensitive to changes in the spectrum continuum, and the errors we estimate from the covariance matrix alone might be too small.

We did not find any signs of a systematic bias in any of the analysis we carried out.

4.4.2 Real fits

In this section we discuss the quality of the fits to SDSS galaxies obtained with VESPA.

As explained in Section 2, VESPA finds the best fit solution in a χ^2 sense for a given parametrization, which is self-regulated in order to not allow an excessive number of fitting parameters. We have shown that this self-regularization gives a better solution in parameter space (Figure 4.15), despite often not allowing the parametrization which would yield the best fit in data space (Figure 4.7). However, our aim is still to find

a solution which gives a good fit to the real spectrum. Figure 4.19 shows the 1-point distribution of reduced values of reduced χ^2 for 1 plate of galaxies. This distribution peaks at around $\chi^2_{reduced} = 1.3$, and figure 4.20 shows a fit to one of the galaxies with a typical value of goodness of fit.

It is worth noting that the majority of the fits which are most pleasing to the eye, correspond to the ones with a high signal to noise ratio and high value of reduced χ^2. One would expect the best fits to come from the galaxies with the best signal. However, we believe the fact that they do not is not a limitation of the method, but a limitation of the modelling. There are a number of reasons why VESPA would be unable to produce very good fits to the SDSS data. One is the adoption of a single velocity dispersion (170 kms^{-1}) which could easily be improved upon at the expense of CPU time. However, the dominant reason is likely to be lack of accuracy in stellar and dust modelling - whereas BC03 models can and do reproduce a lot of the observed features, it is also well known that this success is limited as there are certain spectral features not yet accurately modelled, or even modelled at all. There are similar deficiencies in dust models and dust extinction curves. The effect of the choice of modelling should not be overlooked, and we refer the reader to a discussion in Section 4.5 of Panter et al. (2007), where these issues are discussed.

4.4.3 VESPA and MOPED

In this section we take the opportunity to compare the results from VESPA and MOPED, obtained from the same sample of galaxies. The VESPA solutions used here are obtained with a one-parameter dust model, to allow a more fair comparison between the two methods. Both methods make similar assumptions regarding stellar models, but MOPED uses an LMC (Gordon et al. 2003) dust extinction curve, and single screen modelling for all optical depths.

Our sample consist of two plates from the SDSS DR3 (Abazajian et al. 2005) (plates 0288 and 0444), from which we analyse a total of 821 galaxies. We are mainly interested in comparing the results in a global sense. MOPED in its standard configuration attempts to recover 23 parameters (11 star formation fractions, 11 metallicities and 1 dust parameter), so we might expect considerable degeneracies. Indeed, in the past the authors of MOPED have cautioned against using it to interpret individual galaxy

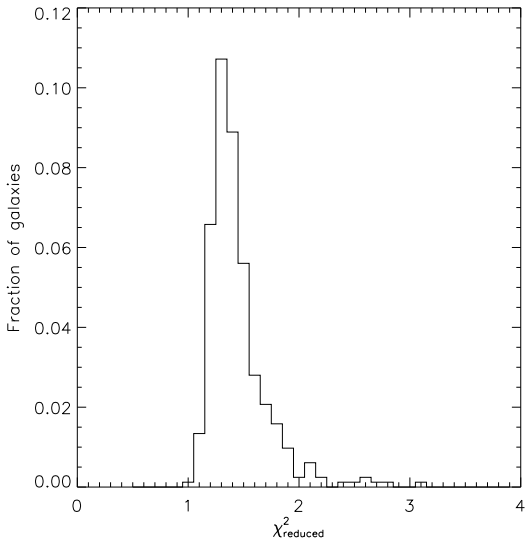

Figure 4.19: The distribution of reduced values of χ^2 for a sample of 360 galaxies analysed by VESPA.

spectra too precisely. We have observed degeneracies between adjacent bins in MOPED, but on the other hand a typical MOPED solution has many star formation fractions which are essentially zero, so the number of significant contributions is always much less than 23.

Figure 4.21 shows the recovered average star formation history for the 821 galaxies using both methods. In the case of VESPA, solutions parametrized by low-resolution bins had to be re-parametrized in high-resolution bins, so that a common grid across all galaxies could be used. This was done using the weights given by (4.22). The lines show a remarkably good agreement between the two methods.

Having recovered a star formation history for each galaxy, one can then estimate the stellar mass of a galaxy. We calculated this quantity for all galaxies using the solutions from both methods, and with similar assumptions regarding cosmological parameters and fibre-size corrections. Explicitly, we have done the following:

1. We converted from flux to luminosity assuming the set of cosmological parameters given by Spergel et al. (2003).

2. We recovered the initial mass in each age bin using each method.

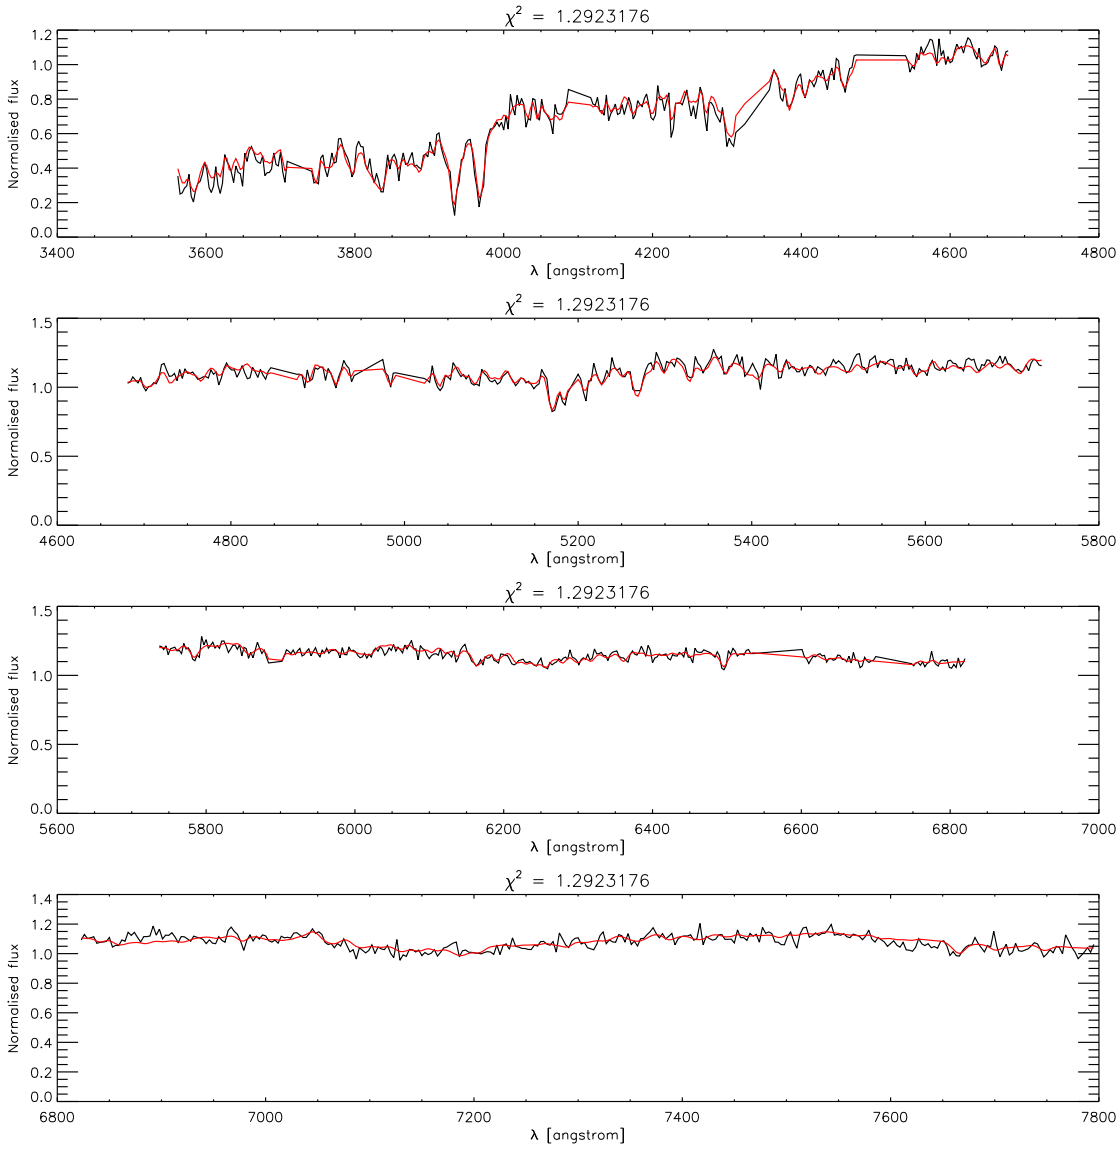

Figure 4.20: Typical fit to a galaxy from the SDSS. The dark line is the real data (arbitrary normalisation), and the red line is VESPA's fit to the data.

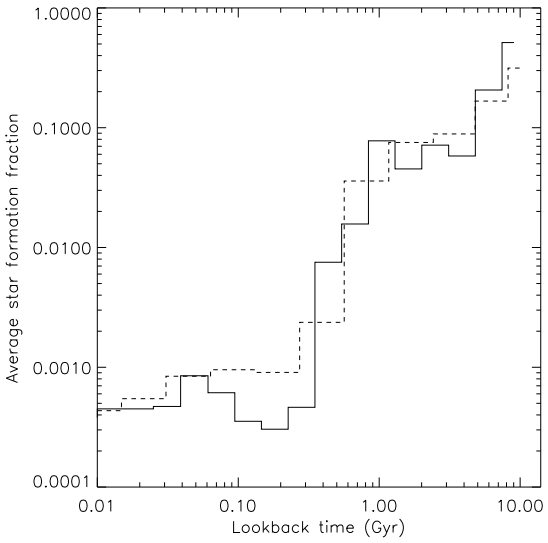

Figure 4.21: The recovered average star formation history for the 821 galaxies as recovered by VESPA (solid line) and MOPED (dashed line). Both were initially normalised such that the sum over all bins is 1, and the MOPED line was then adjusted by 11/16 to account for the different number of bins used in each method, to facilitate direct comparison.

3. We calculated the remaining present-day mass for each population after recycling processes. This information is supplied by the synthetic stellar models, as a function of age and metallicity.

4. We summed this across all bins to calculate the total stellar present-day mass in the fibre aperture, M.

5. We corrected for the aperture size by scaling up the mass to $M_{stellar}$ using the petrosian and fibre magnitudes in the z-band, $M_p(z)$ and $M_f(z)$, with: $M_{stellar} = M \times 10^{0.4[M_p(z)-M_f(z)]}$

Figure 4.22 shows the recovered galaxy masses as recovered from MOPED and from VESPA. We see considerable agreement between VESPA and MOPED. Over 75 per cent of galaxies have $0.5 \leq M_{VESPA}/M_{MOPED} \leq 1.5$. There is a tail of around 10 per cent of galaxies where VESPA recovers 2 to 4 times the mass recovered by MOPED. The main reason for this difference is in the dust model used - we find a correlation between dust extinction and the ratio of the two mass estimates. This again reflects the fact that total stellar mass estimates are highly sensitive to changes in the spectrum continuum (see also section 4.4.1).

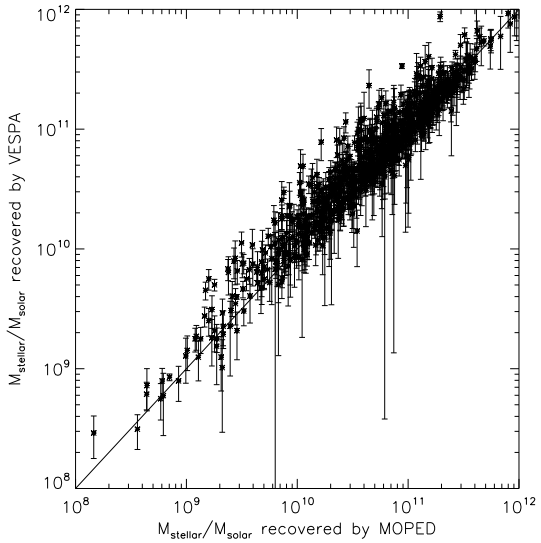

Figure 4.22: Galaxy stellar mass (in units of solar masses) as recovered by VESPA and MOPED for a sub-sample of 821 SDSS galaxies. The small percentage of galaxies with significantly larger VESPA masses have large extinction. The difference is accounted for by the fact that MOPED and VESPA use different dust models.

Our sub-sample of 821 includes galaxies with a wide range of signal-to-noise ratios, star formation histories and even wavelength range (mainly due to each galaxy having different masks applied to it, according to the quality of the spectroscopic data). Figure 4.23 shows the number of recovered non-zero parameters in the sample, using VESPA. As an average, it falls below the synthetic examples studied in Section 3. This is not surprising, though, as each galaxy will have a unique and somewhat random combination of characteristics which will lead to a different number of parameters being recovered. The total combination of these sets of characteristics would be impossible to investigate using the empirical method described in Section 3, and here lies the advantage of VESPA of dynamically adapting to each individual case. Also important to note is the fact that the wavelength coverage is normally not continuous in an SDSS galaxy, due to masked regions. This was not modelled in Section 3, and is likely to further reduce the number of recovered parameters in any given case.

Perhaps more useful is to translate this number into a number of recovered significant stellar populations for each galaxy. We define a significant component as a stellar

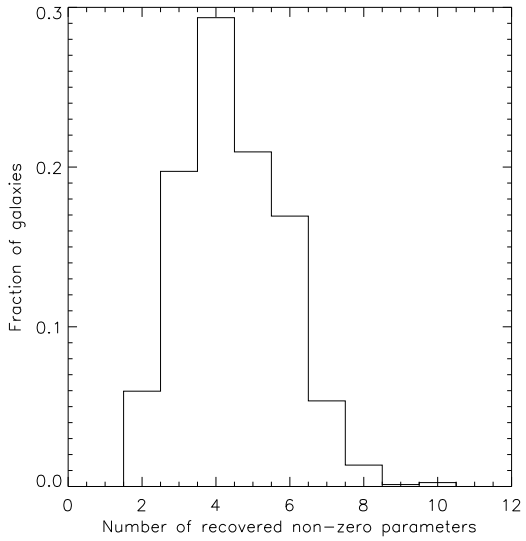

Figure 4.23: Number of non-zero parameters in solutions recovered from 821 SDSS galaxies with VESPA. Please note that these correspond to the *total* number of non-zero components in the solution vector $\mathbf{c}_\kappa$ and not to the number of recovered stellar populations. For information about the number of recovered populations see Figure 14.

population which contributes 5 per cent or more to the total flux. Figure 4.24 shows the distribution of the number of significant components for our sub-sample of galaxies, as recovered by MOPED and VESPA. It is interesting to note that both methods recover on average a similar amount of components, even though MOPED has no explicit self-regularization mechanism, as VESPA clearly does.

4.5 Final remarks on VESPA

We have developed a new method to recover star formation and metallicity histories from integrated galactic spectra - VESPA. Motivated by the current limitations of other methods which aim to do the same, our goal was to develop an algorithm which is robust on a galaxy-by-galaxy basis. VESPA works with a dynamic parametrization of the star formation history, and is able to adapt the number of parameters it attempts to recover from a given galaxy according to its spectrum. In this thesis we tested VESPA against a series of idealised synthetic situations, and against SDSS data by comparing our results with those obtained with the well-established code, MOPED.

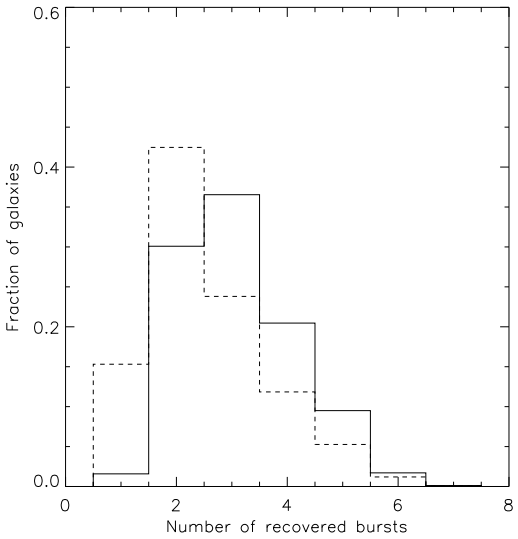

Figure 4.24: The distribution of the total number of recovered stellar populations which contribute 5 per-cent or more to the total flux of the galaxy, as recovered from MOPED (dashed line) and VESPA (solid line).

Using synthetic data we found the quality and resolution of the recovered solutions varied with factors such as type of star formation history, noise in the data and wavelength coverage. In the vast majority of cases, and within the estimated errors and bin-correlations, we observed a reliable reproduction of the input parameters. As the signal-to-noise decreases, it becomes increasingly difficult to recover robust solutions. Whereas our method cannot guarantee a perfect solution, we have shown that the self-regularization we imposed helped obtain a cleaner solution in an overwhelming majority of the cases studied.

On the real data analysis, we have studied possible effects from systematics using duplicate observations of the same set of galaxies, and have also compared VESPA's to MOPED's results obtained using the same data sample. We found that in the majority of cases our results are robust to possible systematics effects, but that in certain cases and particularly when calculating stellar masses, VESPA might underestimate the mass errors. However, we found no systematic bias in any of our tests. We have also shown that VESPA's results are in good agreement with those of MOPED for the same sample of galaxies. VESPA and MOPED are two fundamentally different approaches to the

same problem, and we found good agreement both in a global sense by looking at the average star formation history of the sample, and in an individual basis by looking at the recovered stellar masses of each galaxy. VESPA typically recovered between 2 to 5 stellar populations from the SDSS sample.

VESPA's ability to adapt dynamically to each galaxy and to extract only as much information as the data warrant is a completely new way to tackle the problem of extracting information from galactic spectra. Our claim is that, for the most part, VESPA's results are robust for any given galaxy, but our claim comes with two words of caution. The first one concerns very noisy galaxies - in extreme cases (SNR$\approx$10 or less, at a resolution of 3 Å), it becomes very difficult to extract any meaningful information from the data. This uncertainty is evident in the large error bars and bin-correlations, and the solutions can be essentially unconstrained even at low-resolutions. We are therefore limited when it comes to analysing individual high-noise galaxies, which is the case of many SDSS objects. Our second word of caution concerns the stellar models used to analyse real galaxies - any method can only do as well as the models it bases itself upon. We are limited in our knowledge and ability to reproduce realistic synthetic models of stellar populations, and this is inevitably reflected in the solutions we obtain by using them. On the plus side, VESPA works with any set of synthetic models and can take advantage of improved versions as they are developed.

VESPA is fast enough to use on large spectroscopic samples (a typical SDSS galaxy takes 1 minute on an average workstation), and we have applied VESPA to SDSS's Data Release 5 (DR5), which consists of roughly half a million galaxies. The resulting catalogue and science results are presented in the next two chapters.

Chapter 5

A catalogue of star formation histories

VESPA can be applied to any spectrum, in any wavelength range and at any resolution. We have applied VESPA to SDSS's 5th data release (Adelman-McCarthy et al. 2007), and constructed a public catalogue of star formation histories and other derived products. This Chapter describes the survey, the main galaxy sample which we analysed, and the VESPA catalogue itself. We also present some basic results obtained by exploiting this catalogue, such as a local and evolved galaxy stellar mass function.

5.1 The Sloan Digital Sky Survey

The SDSS is a photometric and spectroscopic survey which images the sky from a dedicated 2.5m telescope in Apache Point, New Mexico. The photometry is done in five bands: u, g, r, i and z corresponding to central effective wavelengths of 3590Å 4810Å 6230Å 7640Å and 9060Å respectively. For details on the hardware, software and data products see York et al. (2000a) and Stoughton et al. (2002). Briefly, the survey is carried out by a mosaic CCD camera (Gunn et al. 1998), two 3 arcsec bre-fed spectrographs, and an auxiliary 0.5m telescope for photometric calibration.

The sky area covered has been incremented since the first area release, and at the time of writing the latest data release (data release 6, Adelman-McCarthy et al. (2008))

covers 9583 square degrees, whereas the spectroscopic survey covers 7425 square degrees - and supplies over 1 million spectra. The key technical aspects of DR5, which is the data release used in this Chapter, are summarised in Table 5.1.

Imaging area	8000 sq. deg
Imaging catalogue	215 million unique objects
Magnitude limits	$22.0(u'), 22.2(g'), 22.2(r'), 21.3(i'), 20.5(z')$
Spectroscopic area	5740 sq. deg.
Wavelength coverage	3800-9200 Å
Resolution	1800
Signal to noise	> 4 per element at g=20.2
Galaxy spectra	674,749

Table 5.1: SDSS's fifth data release.

5.1.1 The main galaxy sample

The main galaxy sample (Strauss et al. 2002b) is a magnitude-limited, high-completeness ($> 99\%$) galaxy sample, selected in the r-band. At the time of the fifth data release, it yields approximately 90 galaxies per square degree with a median redshift $\tilde{z} = 0.11$, totalling over 465,000 galaxies.

Galaxies are selected according to three criteria: a star-galaxy separation test, a cut in the r'-band petrosian magnitude (from here on we will refer to SDSS's r'- as r-band), and a cut in surface brightness. The star-galaxy separation is obtained by comparing the best fit of exponential or de Vaucouleurs model magnitude with the object's point spread function magnitude and requires $r_{PSF} - r_{model} \geq 0.3$. The petrosian radius, θ_p, is defined as the radius at which the local surface brightness in an annulus about θ_p is 1/5 of the mean surface brightness within θ_p. This leads to the definition of petrosian magnitude, r_p, as the flux within an aperture of $2\theta_p$ and is required to obey $15.0 \leq r_p \leq 17.77$. The surface brightness is not calculated within θ_p, but θ_{50} instead - the half-light radius which is defined as the radius which encloses half the petrosian flux - and $\mu_{50} = r_p + 2.5(2\pi\theta_{50}^2)$ must follow $\mu_{50} \leq 24.4$. To avoid very poor quality data we impose a further limit of $\mu_{50} < 23$, and do not analyse galaxies which have all their flux pixels masked.

5.1.2 Spectro-photometric callibrations

In DR5 the spectra are calibrated using standard stars. On each spectroscopic plate 16 standards are observed, mostly sub-dwarf stars. A model spectrum is fitted to each standard, and a flux calibration vector is constructed by calculating the average ratio of each standard's spectrum to its best-fit model. The calibration is done separately in the blue and red channels. The observed spectra in the blue and red arms are then multiplied by the flux calibration method, combined, and re-binned to a constant dispersion.

5.1.3 Galactic extinction

Spectra in DR5 do not include a correction for dust extinction due to our own galaxy. We use the Galactic dust maps by Schlegel et al. (1998) to obtain a value of E(B-V) for each spectroscopic plate. Using the dust extinction curve of O'Donnell (1994) and equation (1.80) we calculate A_λ. We multiply our flux and noise vectors by $10^{0.4A_\lambda}$ (λ in the Earth's frame), giving us an estimate of the un-obscured flux. This assumes a uniform dust screen.

5.1.4 Handling SDSS data

Prior to any analysis, we processed the SDSS spectroscopic data, so as to accomplish the desired spectral resolution and mask out any unwanted signal.

The SDSS data-files supply a mask vector, which flags any potential problems with the measured signal on a pixel-by-pixel basis. We use this mask to remove any unwanted regions and emission lines. In practical terms, we ignore any pixel for which the provided mask value is not zero.

The BC03 synthetic models produce outputs at a resolution of 3Å, which we convolve with a Gaussian velocity dispersion curve with a stellar velocity $\sigma_V = 170\text{km}s^{-1}$, this being a typical value for SDSS galaxies (Panter et al. 2007). We take the models' tabulated wavelength values as a fixed grid and re-bin the SDSS data into this grid, using an inverse-variance weighted average. We compute the new error vector accordingly. Note that the number of SDSS data points averaged into any new bin is not constant, and that the re-binning process is done after we have masked out any unwanted pixels. Additionally to the lines yielded by the mask vector, we mask out the following emission

line regions in every spectrum's rest-frame wavelength range: [5885-5900, 6702-6732, 6716-6746, 6548-6578, 6535-6565, 6569-6599, 4944-4974, 4992-5022, 4846-4876, 4325-4355, 4087-4117, 3711-3741, 7800-11000] Å. These regions were determined by visual inspection of over 1000 galaxy fit residuals (Panter et al. 2007).

These re-binned data- and noise-vectors are essentially the ones we use in our analysis. However, since the linear algebra assumes white-noise, we pre-whiten the data and construct a new flux vector $F'_j = F_j/\sigma_j$, which has unit variance, $\sigma'_j = 1, \forall j$, and a new model matrix $A'_{ij} = A_{ij}/\sigma_j$.

The wavelength vector is shifted to the galaxy's rest frame.

5.2 The catalogue

The catalogue is published as a query-based T-SQL database. The data is organised in tables, which we describe in section 5.2.6. The principal physical properties provided by VESPA are summarised in Table 5.2 .

Symbol	Units	Description
x_α	-	star formation fractions in bin α
u_α	$M_\odot$	stellar mass formed in bin α
m_α	$M_\odot$	recycled stellar mass in bin α
$C_{\alpha\beta}$	$M_\odot^2$	covariance matrix for the star formation fractions
Z_α	-	mass-weighted metallicity for bin α
$C^Z_{\alpha\beta}$	-	covariance matrix for the metallicities
M_*	$M_\odot$	recycled stellar mass in galaxy
M^u_*	$M_\odot$	total stellar mass formed in galaxy
τ_V^{ISM}	-	dust extinction due to the inter-stellar medium for all populations
τ_V^{BC}	-	dust extinction due to the birth cloud for young populations

Table 5.2: Galaxy properties which are derived by VESPA.

Additionally, some derived quantities are calculated, as shown in Table 5.3.

Next we describe the meaning and calculation of each of these quantities.

Symbol	Units	Description
V_{max}	Mpc^3	maximum volume a galaxy could be observed in, given its petrosian r magnitude, surface brightness and star formation history
$M_*(z)$	$M_\odot$	stellar mass in a given galaxy at a redshift $z > z_{gal}$, where z_{gal} is the redshift of the galaxy, and assuming a history with no mergers

Table 5.3: Additional properties derived from VESPA's output.

5.2.1 Masses and mass fractions

VESPA recovers stellar mass fractions per age bin. We transform these mass fractions into absolute masses as follows.

We measure a flux, $F = \frac{Energy}{\Delta t \Delta\lambda \Delta A}$, whereas the models are given in units of solar luminosity, $L = \frac{Energy}{\Delta t \Delta\lambda}$. In relation to the source, we sit on the surface of a sphere of proper radius $R_0 S_k$ and surface area $4\pi R_0^2 S_k^2$ but the relation between flux and luminosity is not as simple as $F_\lambda = \frac{L_\lambda}{4\pi R_0^2 S_k^2}$ due to the effects of the expansion of the universe. We have time dilation: $\Delta t_{obs} \rightarrow \Delta t_{em}(1+z)$; photon energies being redshifted: $E_{obs} \rightarrow E_{em}/(1+z)$, and bandwidth in wavelength being increased: $\Delta\lambda_{obs} \rightarrow \Delta\lambda_{em}(1+z)$. Combining all of these effects we get

$$F_\lambda = \frac{L_{\lambda/(1+z)}}{4\pi D_L^2 (1+z)} \tag{5.1}$$

where D_L is the luminosity distance, as defined by equation (1.23).

In practical terms we do the following. Recall equation (4.2):

$$F_j = \sum_\alpha x_\alpha G_{\alpha j} \tag{5.2}$$

where F_j has been shifted to the rest frame of the source. Both the observed flux and the models are normalised to numbers of a few orders of magnitude, to ensure the computational calculations are not affected by very large (or very small) numbers. We write

$$\frac{F_j}{\langle F_j \rangle} = \sum_\alpha x_\alpha \frac{G_{\alpha j}}{\langle G_{00} \rangle}. \tag{5.3}$$

The luminosity of a galaxy, written in terms of our models, is simply $L_j = \sum_\alpha u_\alpha G_{\alpha j}$, where u_α is the stellar mass formed at age α. From (5.1):

$$F_j = \frac{\sum_\alpha u_\alpha G_{\alpha j}}{4\pi D_L^2 (1+z)}. \tag{5.4}$$

For any given age bin α, combining equations (5.3) and (5.4) gives the mass formed in each bin α in units of solar masses:

$$u_\alpha = x_\alpha \frac{\langle F_j \rangle}{\langle G_{00} \rangle} 4\pi D_L^2 (1+z). \tag{5.5}$$

We distinguish between the stellar mass ever formed in a galaxy, and the stellar mass remaining in a galaxy today:

$$M(t) = \int_0^t \psi(t') \left(1 - R(t-t')\right) dt' \tag{5.6}$$

where $R(t-t')$ is the fraction of stellar mass lost to the ISM at time t, by a stellar population aged t' and $\psi(t')$ is the star formation rate at age t'. In practical terms we calculate the following:

$$m_\alpha = u_\alpha R_\alpha \tag{5.7}$$

$$M^u_{*,fibre} = \sum_\alpha u_\alpha \tag{5.8}$$

$$M_{*,fibre} = \sum_\alpha m_\alpha \tag{5.9}$$

where R_α, known as the *recycling fraction* is given by the BC03 models, for each of the metallicities.

Finally, we correct for the fact that the fibre has an aperture of 3 arcseconds, which means we do not typically observe the entirety of the galaxy. We use the observed fibre and petrosian magnitudes in the z-band to scale up the stellar mass as

$$M_* = M_{*,fibre} \times 10^{0.4(z_p - f z_p)} \tag{5.10}$$

where z_p and $f z_p$ are the petrosian and fibre magnitudes in the z-band respectively. This scaling assumes that the parts of the galaxy which do not fall under the fibre's aperture have an identical star formation history as that observed. For an ensemble of galaxies fibre aperture corrections are not important, in the sense that the mean colour from the fibre is the same as the mean colour from the photometry (Glazebrook et al.

2003a). However, one should keep in mind that they remain important for individual galaxies, or at very low redshift.

5.2.2 Error estimates

As mentioned in section 4.2.5, we estimate the effect of the noise in our recovered solutions by explicitly calculating a number of noisy realisations in order to estimate the covariance matrices in (4.23) and (4.24). For convenience, we also define a covariance matrix of the unrecycled mass per bin as (using equation 5.5)

$$C_{\alpha\beta}(u) = C_{\alpha\beta}(x) \left(\frac{\langle F_j \rangle}{\langle G_{00} \rangle} 4\pi D_L^2 (1+z) \right)^2 \tag{5.11}$$

from which we can estimate the error in the unscaled mass formed in bin α as $\sigma_{u(\alpha)} = \sqrt{C_{\alpha\alpha}(u)}$. This ignores the uncertainty in the estimation of the redshift, which is of little significance compared to the variance of u_α across realisations. From $\sigma_{u(\alpha)}$ we can calculate errors in $\sigma_{m(\alpha)}$ by multiplying by the corresponding recycling fraction R_α. The error on the metallicity is simply $\sqrt{C_{\alpha\alpha}(Z)}$.

We use the full covariance matrix to estimate the statistical error on the total stellar mass, M_*:

$$\sigma^2(M_*) = \sigma^2(M_{*,fibre})\gamma^2 + M_{*,fiber}^2 \sigma^2(\gamma) \tag{5.12}$$

where γ is the conversion factor between fibre and galaxy mass of equation (5.10), and $\sigma(\gamma)$ is the error associated with this factor, calculated using the errors in the petrosian and fibre z-band magnitudes. $\sigma^2(M_{*,fibre})$ is estimated from the unrecycled mass covariance matrix and the total recycling fraction of the galaxy, R:

$$\sigma^2(M_{*,fibre}) = \sum_{\alpha,\beta} C_{\alpha\beta}(u) R^2. \tag{5.13}$$

This assumes that there is no error in the recycling fractions - i.e., that we know the SFH exactly.

5.2.3 Mass and metallicity per age bin

The history of each individual galaxy is likely to be parametrized by a combination of high and low resolution age bins. One must be careful to interpret the masses and metallicities associated with low-resolution bins. We recall that in these cases the mass

recovered should be interpreted as the total mass recovered in the bin, but we have little information of when, within the bin, it formed. Similarly, we should interpret the metallicity values recovered as a mass-weighted metallicity for the whole duration of the bin.

5.2.4 Dust

We recover two values of dust extinction, one associated with young stars (τ_V^{BC} - applied to stellar populations younger than 0.03 Gyrs) and one associated with the whole galaxy (τ_V^{ISM} - applied to all stellar populations), following the mixed slab model of Charlot and Fall (2000). These values relate to a dust extinction curve that follows

$$\tau_\lambda = \tau_V \left(\frac{\lambda}{5500\mathring{A}} \right)^{-0.7}. \tag{5.14}$$

A comparison of this curve with the estimated extinction curve of the LMC by Gordon et al. (2003) can be seen in Figure 1.4. In this figure both curves are calibrated such that $\tau_{5550\mathring{A}} = 1$. Uncalibrated optical depth values are related as $\tau^{LMC}_{5550\mathring{A}} = \tau_{5550\mathring{A}}/3.37$.

5.2.5 V_{max}

The main galaxy sample is a magnitude-limited sample, which means that in any given volume we are only observing all of the objects up to a certain magnitude, and will miss fainter objects. If we want to calculate densities of quantities over the entire volume of the survey we must correct for this incompleteness. A common method is the so-called V_{max} method, which calculates the maximum comoving volume in which a given galaxy could be observed. We can then estimate the comoving density of the quantity F by

$$n_F = \sum_i \frac{F_i}{V_{max,i}}. \tag{5.15}$$

A faint galaxy will have a small V_{max}, and a larger weight. Conversely, a very bright galaxy will be observed wherever it is within the survey volume, and be given a weight of one, in units of the volume of the survey. This assumes that the galaxies are uniformly distributed across the sky, which even though it is clearly not true is an acceptable approximation for a volume as large as the one we are dealing with.

The maximum volume is then calculated as

$$V_{max} = \int_{z_{min}}^{z_{max}} \frac{dV}{dz} dz \tag{5.16}$$

where $z_{max} = min\{z^{cut}_{max}, z(r_p, \mu_{50})\}$ and $z_{min} = max\{z^{cut}_{min}, z(r_p, \mu_{50})\}$, where $z^{cut}_{max}, z^{cut}_{min}$ are the redshift limits of the survey or sample selection and $r_p(z)$ and $\mu_{50}(z)$ are calculated as a function of a galaxy's observed r-band and surface brightness magnitudes, and its star formation history.

The r-band magnitude of a galaxy, as a function of redshift, changes due to three main reasons: firstly simply because of distance; secondly because the observed band shifts along the galaxy's spectrum with changing redshift (the so called K-correction), and thirdly because galaxies are evolving objects, whose intrinsic brightness changes across their lifetimes (the E-correction). We can write

$$r_p(z) = M_p + 5\log\frac{D_L(z)}{10pc} + C(z) \tag{5.17}$$

where the second term accounts for distance effects, and the third term for the K- and E-corrections. $C(z)$ is the magnitude difference arising from the ratio of the emitted and observed r-band luminosities:

$$C(z) = -2.5\log\left(\frac{\int_{-\infty}^{\infty} L^{ev}_{\lambda/(1+z)} T_\lambda \frac{d\lambda}{1+z}}{\int_{-\infty}^{\infty} L^{em}_{\lambda/(1+z_{gal})} T_\lambda \frac{d\lambda}{1+z_{gal}}}\right) \tag{5.18}$$

where L^{ev} is the evolved luminosity of a galaxy at a redshift z, L^{em} is the luminosity emitted at $z = z_{gal}$, T_λ is SDSS's r-band filter response and λ is in the Earth's frame. If at any given redshift the luminous output of the galaxy is greater than the emitted at z_{gal} then the magnitude correction is negative (brighter), as expected. For $z < z_{gal}$ we have no information on L^{em}. In this case we take the flux to be that observed at z_{gal} and calculate a K-correction only.

With VESPA's star formation histories, one could in principle reconstruct the flux L^{ev} at any redshift $z > z_{gal}$. However, there are computational implications because of the way the high- and low-resolution bins are constructed - we would have to reconstruct a whole new set of bins at ages dictated both by the redshift of the galaxy and the redshift at which we would like to calculate the correction. This is too computationally expensive, so we have approximated the problem as follows.

For each tabulated metallicity, we estimate the r-band flux as a function of age for a stellar population of $1M_\odot$ at $t = 0$, and construct the vector

$$l(t, Z, z) = \int_{-\infty}^{\infty} L_{\lambda/(1+z)}(t, Z) T_\lambda \frac{d\lambda}{1+z} \tag{5.19}$$

which corresponds to the observed r-band luminosity of a population of stars aged t, metallicity Z, at a redshift z. We calculate this quantity for the five tabulated metallicities and 16 ages corresponding to the high resolution bins detailed in section 4.2.3. For each galaxy, we then simplify its star formation history to a series of bursts which occur at the center of each of its bins. For each bin α we calculate $l_\alpha(t_\alpha, Z_\alpha, z)$ by linearly interpolating the quantity pre-calculated by equation (5.19) over the correct range in age, metallicity and redshift. The age of each burst, t_α, depends on the redshift z and is shifted for each new evaluation at a different redshift. When a population has an age t_α which is less than zero its contribution is set to zero.

We then estimate the total r-band evolved luminosity of any galaxy at a redshift z by adding up the contributions from all its stellar populations:

$$l^{ev}(z) = \sum_\alpha f(t_\alpha, Z_\alpha, z) x_\alpha \tag{5.20}$$

and approximate equation (5.18) by

$$C(z) \approx \frac{l^{em}(z_{gal})}{l^{ev}(z)} \tag{5.21}$$

where l^{em} is simply $l^{ev}(z)$ evaluated at z_{gal}.

Changes in surface brightness can be calculated by the changes in magnitude:

$$\mu_{50}(z) = r_p(z) + 2.5 \log_{10}(2\pi\theta_{50}^2(z)). \tag{5.22}$$

If we assume that the physical size of the galaxy does not change with redshift, then changes in θ_{50} are given by the ratio of the angular-diameter distances. This is a good approximation for the local Universe. We write $\theta_{50}(z) = \theta_{50}(z_{gal}) \frac{D_A(z_{gal})}{D_A(z)}$.

This enables us to track the evolution of observed magnitude and surface brightness for each galaxy as a function of redshift. To calculate V_{max} we need to estimate the total volume in which a given galaxy would have been observed. We use equations (5.17) and (5.22) to estimate where any given galaxy drops in and out of the survey, using the survey selection limits of section 5.1.1 and integrate equation (5.16) over the relevant values of z_{min} and z_{max}.

5.2.6 The database

The catalogue has been published as a relational T-SQL database, which provides maximum flexibility and interoperability with other databases. The database is split into a number of tables, each with a number of fields (columns), which can be accessed and queried via SQL.

We have split the database into four tables: GalProp includes results relative to the galaxy as a whole, BinProp has results which refer to each bin individually, DustProp holds dust information and BinID identifies each VESPA low- or high-resolution bin. Tables 5.4, 5.5, 5.6 and 5.7 detail each of the fields included in the tables. Galaxies are identified by an identifier which is common with SDSS - specObjID. This means that object properties which are already included in the SDSS need not be included in this database, as cross-matching can be done in a straightforward way using specObjID.

Field	Units	Description
specObjID		Matches SDSS's specObjID and uniquely identifies a spectrum.
m_stellar	$M_\odot$	M_* - equation (5.10).
m_stellar_error	$M_\odot$	$\sigma(M_*)$ - equation (5.12).
t_lb	Gyr	Lookback time of galaxy.
chi2		χ^2 of the spectral fit - equation (4.12).
SNR		Signal to noise ratio of the used (unmasked) spectrum, at 3Å resolution.
nbins		Number of recovered bins in a galaxy.
Vmax	Mpc^3	V_{max} as defined by equation (5.16).

Table 5.4: GalProp

5.3 Basic results

The catalogue described in the previous section can be explored in a variety of ways. Here we describe an investigation into the assembly of stellar mass in galaxies. We use both the total stellar mass of each galaxy, alongside its star formation history to calculate a galaxy stellar mass function of the local Universe, a present-day stellar mass density, and the evolution of both of these quantities as a function of redshift.

Field	Units	Description
specObjID		Matches SDSS's specObjID and uniquely identifies a spectrum.
binID		Bin identifier as given by Figure 4.6.
mass	$M_\odot$	Mass formed in bin - equation (5.5).
mass_error	$M_\odot$	$\sigma_{u(\alpha)}$ - as derived from equation (5.11).
SFR	$M_\odot \mathrm{Gyr}^{-1}$	Star formation rate in bin.
Z		Metallicity in bin.
Z_error		$\sigma_{Z(\alpha)}$ as derived from equation (4.24).

Table 5.5: BinProp

Field	Units	Description
specObjID		Matches SDSS's specObjID and uniquely identifies a spectrum.
dustID		Dust identifier: 1 for τ_V^{BC} and 2 for τ_V^{ISM}.
dustVal		Either τ_V^{BC} or τ_V^{ISM}, according to dustID.

Table 5.6: DustProp

Field	Units	Description
binID		Bin identifier, as given by Figure 4.6.
ageStart	Gyr	Age of the young boundary of the bin.
ageEnd	Gyr	Age of the old boundary of the bin.
width	Gyr	Width of the bin in Gyrs.

Table 5.7: BinID

5.3.1 The observed galaxy stellar mass function

The galaxy stellar mass function (GSMF) gives the comoving number density of galaxies per unit of logarithmic stellar mass, $\phi(M)d\log M = n(M)d\log M$.

It provides a description of how stellar mass is distributed within galaxies, and is one of the fundamental observables which drives models for galaxy formation and evolution. As introduced in section 1.4.5, the current hierarchical model for galaxy formation provides a good explanation for the galaxy stellar mass function we observe, within the context of cold dark matter.

Figure 5.1 shows the galaxy stellar mass function estimated from SDSS's DR5 main galaxy sample. The error bars in each mass bin are estimated using the jackknife method, as described in Wall and Jenkins (2003), pp 132. Briefly, jackknife methods estimate the error on a parameter, θ, estimated from a dataset using the data set alone. This works by creating suitable sub-samples of data points and estimating θ_s from each subsample, as well as $\sigma^2(\theta_s)$ - the variance of θ_s across these subsamples. This can then be related to $\sigma^2(\theta)$. The GSMF is normally fitted with a Schechter function of the form

$$\phi(M)dM = \phi^* \left(\frac{M}{M^*}\right)^{\alpha} \exp\left(\frac{-M}{M_*}\right) d\left(\frac{M}{M_*}\right) \tag{5.23}$$

which we write as

$$\phi(M)d\log M = \phi^{*'} \left(\frac{M}{M^*}\right)^{\alpha'} \exp\left(\frac{-M}{M_*}\right) d\log M \tag{5.24}$$

where $\alpha' = \alpha+1$ and $\phi^{*'} = \phi^* \ln(10)$. The best fitting parameters, for the mass range $9 < \log M/M_\odot < 12$, are: $\phi^{*'} = 0.00432 \pm 0.000186$, $M_* = 1.668 \times 10^{11} \pm 2.247 \times 10^9$ and $\alpha' = -0.255 \pm 0.0147$.

A comparison with other local estimates of the galaxy stellar mass function, both from SDSS and other datasets, is seen in Figure 5.2. Cole et al. (2001) use the 2dF and 2MASS surveys to estimate stellar masses using broadband colours and SSP fits. Bell et al. (2003) cross-match SDSS's early data release with 2MASS using a similar method. Panter et al. (2007) use full spectral fitting to obtain stellar mass estimates for the SDSS-DR3 galaxies. All results are in good agreement within the mass ranges

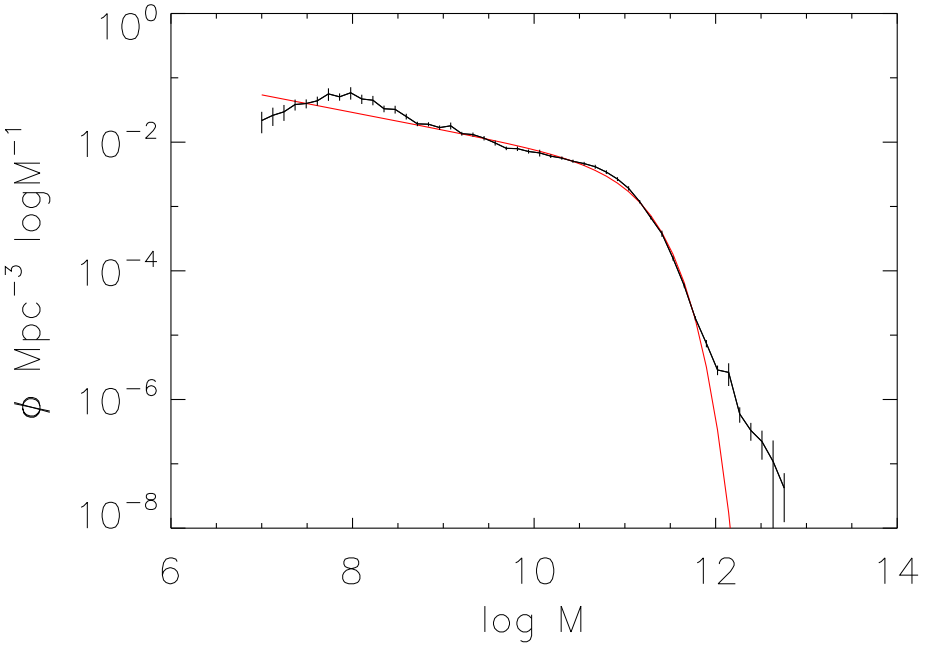

Figure 5.1: The measured galaxy stellar mass function of the SDSS's main galaxy sample for $0.005 < z < 0.35$. The error bars are jackknife estimates and the red line is the best fit Schechter function - see text for details.

in which they overlap, after correcting for differences in the IMF.

We see some evidence for an upturn on the low mass (faint) end. A similar upturn has been seen, both in the luminosity function (e.g. Blanton et al. 2005) and also in the mass function (e.g. Baldry et al. 2008). Baldry et al. (2008) also looked at how the slope of this upturn changes with environment, finding that it is significantly steeper in dense environments. The type of galaxies which populate this mass range is also correlated with environment, suggesting that the physical mechanisms behind the steepening of the slope at low masses are different in clusters and in the field. The downturn for $M/M_{\odot} < 8$ is most likely due to incompleteness in the sample at these masses.

At the high mass end, for $M_* > 10^{12} M_{\odot}$, we see a clear deviation from the Schechter fit or the GSMF of Panter et al. (2007), which is the only one to probe equally high masses. One should be careful not to over-interpret this excess - the number of galaxies

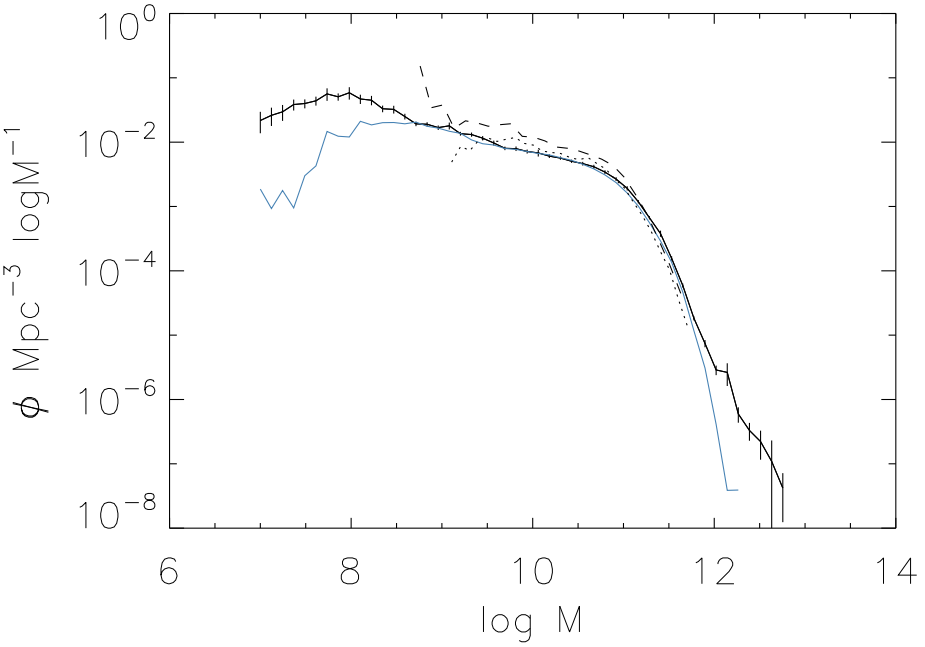

Figure 5.2: The same curve shown in Figure 5.1, this time plotted alongside other estimates of the local GSMF from the literature. Cole et al. (2001) in the solid line, Bell et al. (2003) in the dotted line and Panter et al. (2007) in the solid blue line, all corrected for changes in the assumed IMF if necessary.

per bin decreases steeply from 72 galaxies at $M_* = 10^{12.1} M_\odot$ to one galaxy at $M_* = 10^{12.8} M_\odot$.

5.3.2 The evolution of the GSMF

We have already seen how the observed current population of galaxies can be explained by invoking a hierarchical scenario, in which small galaxies merge to form larger galaxies. This means that the galaxy stellar mass function must have looked very different in past times - we expect both a change in the total number of galaxies in a comoving volume (due to merging) and a qualitative change in the shape of the GSMF due to evolution in stellar mass of individual galaxies, due to star formation and recycling processes. Knowing how the GSMF changes as a function of redshift tells us about how galaxies assemble their stellar mass, which in turn is a fundamental constraint for galaxy formation and evolution models.

The observed evolution of the GSMF

Traditionally one looks at galaxies at different redshift ranges. Figure 5.3 shows the GSMF calculated in four redshift bins, with upper boundaries at 0.05, 0.15, 0.25 and 0.35. Within each redshift slice, a simple $1/V_{max}$ weighting will fail to correct for the galaxies which are too faint to be observed at that redshift at all. This starts at $M_* \approx 10^{10} M_\odot$ for $0.05 < z < 0.15$ and goes up to $M_* \approx 10^{11.5} M_\odot$ at $0.25 < z < 0.25$. For this reason, we only plot the GSMF above these mass limits for each redshift slice. We can note however that the information from the lowest redshift bin can be used to correct for this incompleteness at higher redshifts, and we intend to do this in the future. At low redshift ($z < 0.05$) SDSS's selection function descriminates against high-mass galaxies for fear of shredding. This is an effect caused by large galaxies at low redshift having an apparent angular size which is too large and leads to them being treated as many smaller-sized objects.

The agreement between the lines at different redshifts indicates no evidence for evolution in the galaxy stellar mass function between $z = 0.05$ and $z = 0.35$ for masses higher than $10^{11.5} M_\odot$.

The inferred evolution of the galaxy stellar mass function

Given the information provided by VESPA, we can also estimate the total stellar mass content formed up to any given time in any individual galaxy by the means of equation (5.6). This takes into account both the stars formed up to a time t and the amount of material returned to the inter-galactic medium, and is an estimate of the current stellar mass of a galaxy at time t. In practice, for any redshift z and a galaxy with redshift z_{gal}, we calculate the corresponding lookback time in the galaxy's frame. We then integrate equation (5.6), taking into account whether the time we are interested in has fallen in a high- or low-resolution age bin, given that each assumes a different star formation rate within it - see equations (4.19) and (4.22). This means that, in principle, we can construct a galaxy mass function as a function of redshift, $\phi(M, z)$.

It is not, however, a true representation of the galaxy mass function at those redshifts. The reason is simple: in doing so we are conserving the total number of galaxies between then and now. i.e., we assume a history with no mergers. This limitation arises from the fact that VESPA tells us when stars formed and in which galaxies they

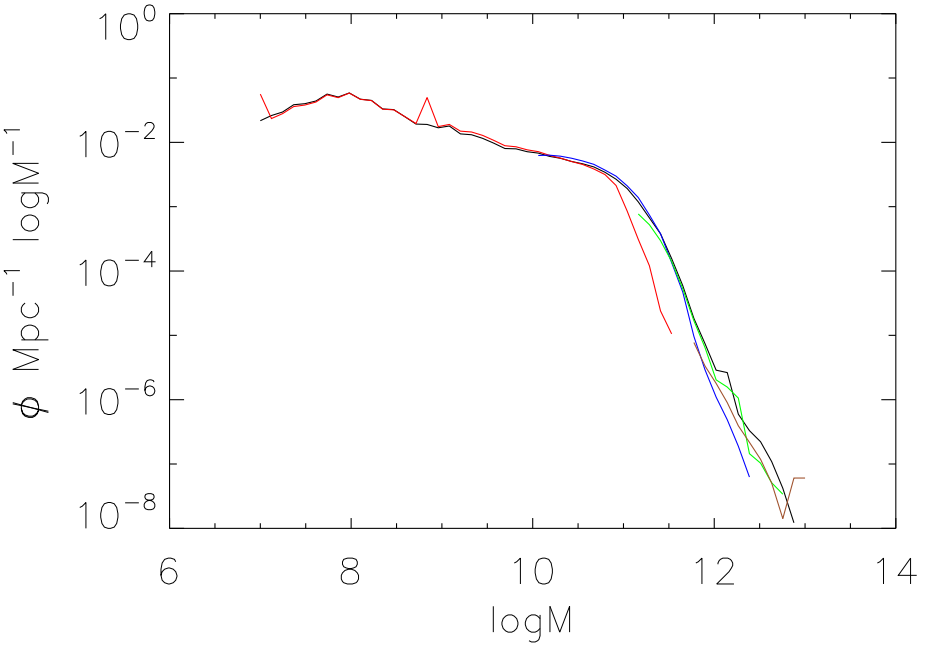

Figure 5.3: The SDSS DR5 galaxy stellar mass function split by redshift bins. The black line shows the mass function for all the galaxies, the red line for galaxies in $0.005 < z < 0.05$, the blue line for galaxies in $0.05 < z < 0.15$, the green line for galaxies in $0.15 < z < 0.25$ and finally the brown line for galaxies in $0.25 < z < 0.35$. The incompleteness limits can be seen creeping up in mass as we go further in redshift (see text). At low redshift we see some effect at high masses due to the SDSS's selection function which discriminates against high mass galaxies due to fear of shredding (see text). We see no sign of evolution in the mass range fully covered by all the redshifts.

end up today (or at the redshift of observation), but not in which galaxies they formed.

Even though, at least formally, we cannot draw any definite conclusions from evolutionary changes in the GSMF as calculated above, it is still very instructive to see how large the differences are. The main reason is that even though one cannot interpret any difference with much confidence, if these differences are significant one might hope to compare the inferred GSMF at a redshift z as calculated from VESPA with one at the same redshift, estimated from methods which directly probe the galaxies at that redshift.

The *difference* between the two measurements would provide a simple galaxy-counting mechanism, which corresponds to disentangling galaxy stellar mass growth due to star formation (and recycling) and galaxy stellar mass growth through mergers. The idea of using the galaxy stellar mass function and its derivative with respect to time, together with information on the star formation history, has already been explored in Drory and Alvarez (2008). In this paper, the authors used the data in the Fors Deep Field, which spans a redshift range $0 < z < 5$ and mass range $10^9 < M_*/M_\odot < 10^{12}$ to estimate how much of galaxy stellar growth is due to mergers, as a function of galaxy stellar mass and redshift. They are mainly limited by their data, and the lack of star formation histories per galaxy. Instead, the authors fit the measured average star formation rate as a function of mass with an analytical expression, which they then integrate to obtain average star formation histories as a function of mass. We therefore expect a dataset as the one provided by VESPA to be a vast improvement on the amount and reliability of information one can extract concerning the merger rate using this kind of analysis.

Figure 5.4 shows the inferred GSMF at the redshifts of 0.5 and 0.9. On the face of it we see no evolution in galaxies with $M > 10^{11} M_\odot$, and see some evidence of low mass galaxies turning into intermediate mass galaxies between the two redshifts. We re-iterate that this keeps track of the changes in mass due to star formation and recycling only - in reality we do not know in how many pieces each of our present-day galaxies are at these redshifts.

Even though the changes are significant within our own sample, we find they are overshadowed when compared to measurements from different samples and their un-

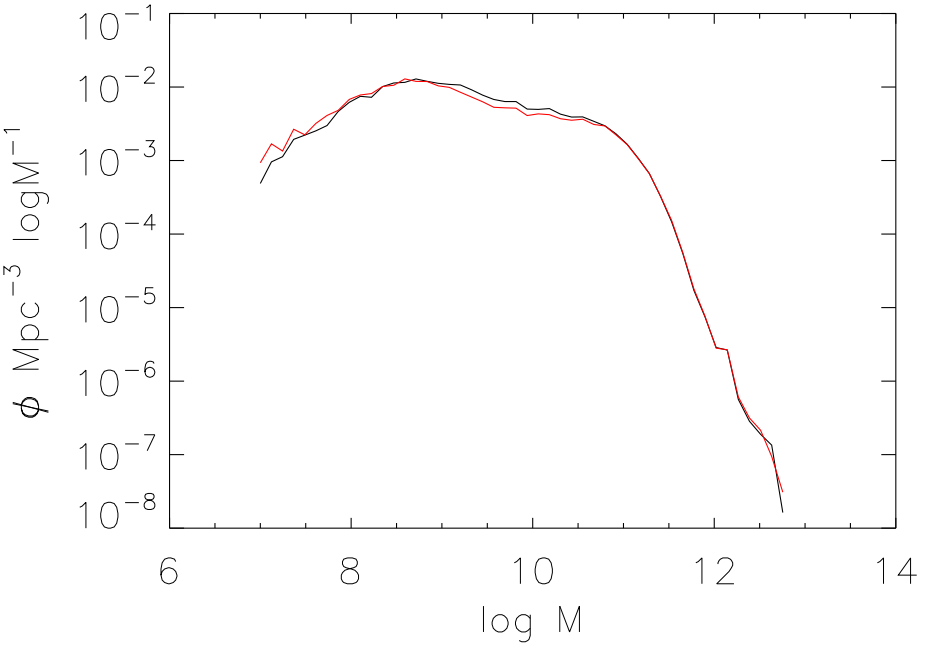

Figure 5.4: The inferred evolution of the GSMF between the redshift of 0.5 (black line) and 0.9 (red line), constructed by integrating the star formation history of each galaxy up to these redshifts.

certainties, as shown in Figure 5.5. This makes it difficult to interpret any difference between the observed and the inferred GSMF at any redshift - they are probably mostly due to differences in data analysis methods, and small compared to estimated errors.

A possible way forward is to use the rate of change of each of the measurements separately. As an example, consider the change in the GSMF between $z = 0.5$ and $z = 0.9$ from the COMBO-17 analysis, shown in Figure 5.6. We show the Schechter fits as published in Borch et al. (2006). The evolution in the parameters is small compared with their uncertainties and we advise caution against interpreting this too closely. However, we can observe how the change in the observed GSMF is qualitatively similar to the change in the inferred GSMF at the high mass end - both measurements see no clear sign of evolution. The same can be said for the mass range down to $\log(M/M_{\odot}) \approx 9$, below which we are most likely incomplete - in this case we see tenuous hints of evolution in both cases. If we could quantify these changes robustly, we would be able to get a handle on how much of the change in the observed GSMF was due to something other

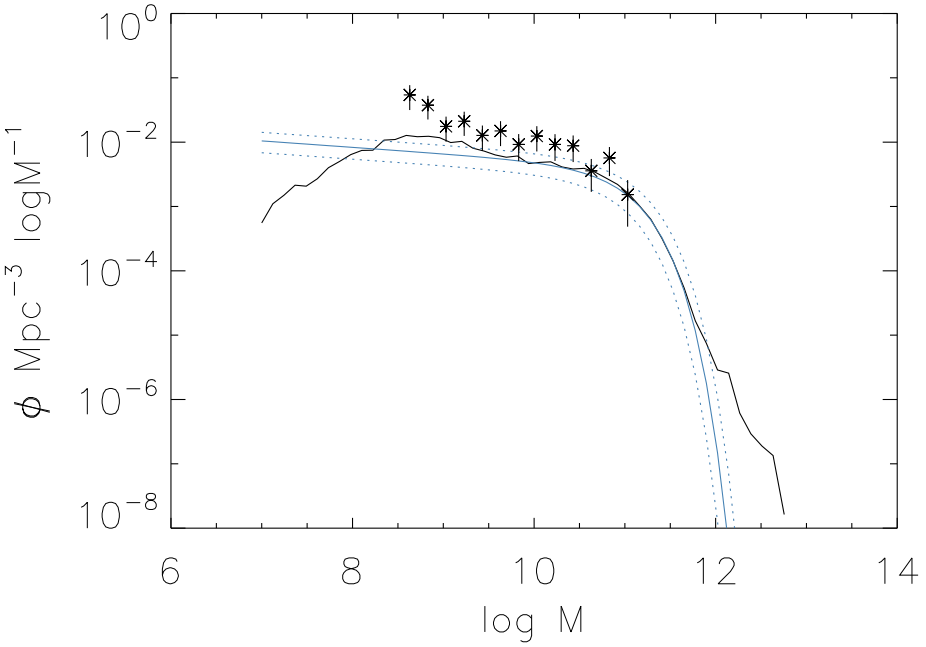

Figure 5.5: The inferred evolution of the GSMF between the redshift of 0.5 as estimated from VESPA on the solid line compared to two estimates of $\phi(M, z = 0.5)$ measured directly at those redshifts: stars from Fontana et al. (2006) and the GOODS-MUSIC survey, and the blue lines are the best-fit Schechter function from Borch et al. (2006) and the COMBO-17 survey (dotted lines are 1-sigma uncertainties). The interpretation of VESPA's curve is different from that of the other two, see text.

than star formation. Unfortunately, the uncertainties are too high for such a study.

Ideally one would like to estimate the inferred and observed GSMF from the same sample or, given the lack of redshift surveys which are complete across a long redshift baseline, to use the exact same analysis methods for spectral or photometric analysis, mass estimates and completeness corrections. We plan to extend this work in this way, taking advantage of future redshift surveys. In the meantime, a public catalogue of $M_*(z)$ for SDSS's DR5 galaxies is available to the community.

5.3.3 Stellar mass density

We can integrate the GSMF over the available mass range in order to estimate a stellar mass density:

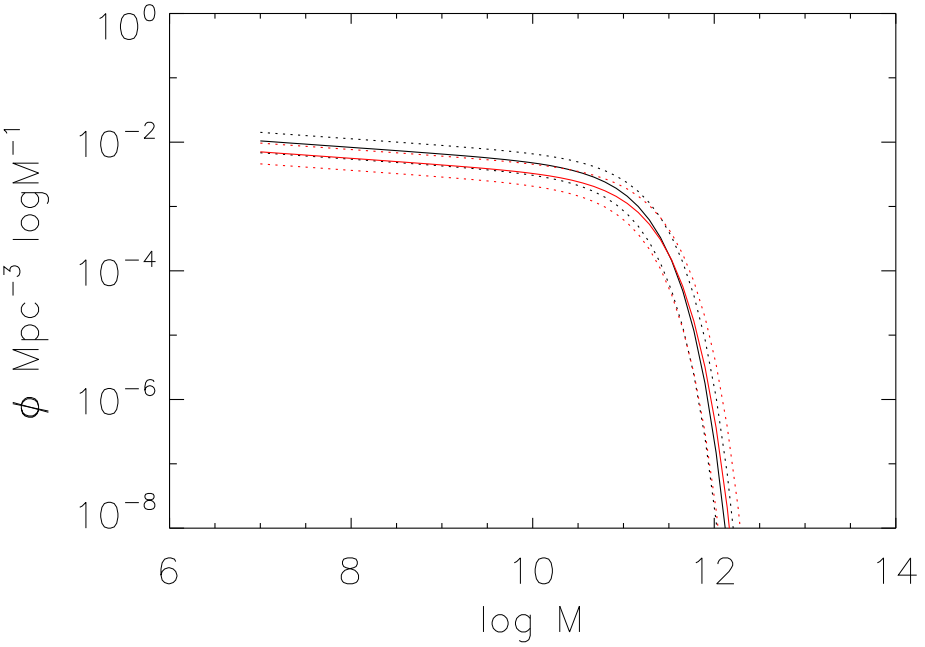

Figure 5.6: The observed evolution of the GSMF between the redshift of 0.5 (black line) and 0.9 (red line), from the COMBO-17 data.

$$\rho_*(z) = \int \phi(M, z) M(z) dM. \tag{5.25}$$

We now lose all sensitivity to how stellar mass is distributed within galaxies, and are simply interested in the total stellar content of the Universe. We describe the stellar mass density in terms of the critical density Ω_*, defined by equation (1.9). Table 5.8 and Figure 5.7 show $\Omega_*(z)$ as calculated from VESPA, integrating masses from $10^7 M_\odot$ to $10^{13} M_\odot$. The points in Figure 5.7 have been converted to a Salpeter IMF to allow comparison with published values.

Overall we find good agreement between the stellar mass density obtained from VESPA and other methods. The large scatter in the black stars of Figure 5.7, corresponding to the coloured points in Figure 4.3, is due to the wide range of data sets and mass estimate methods, but a clear decrease in stellar mass content of the Universe can be seen, as we would expect. VESPA's last point however, at $z = 1.8$, indicates a slight increase in the stellar mass function which is not seen in other measurements. For the

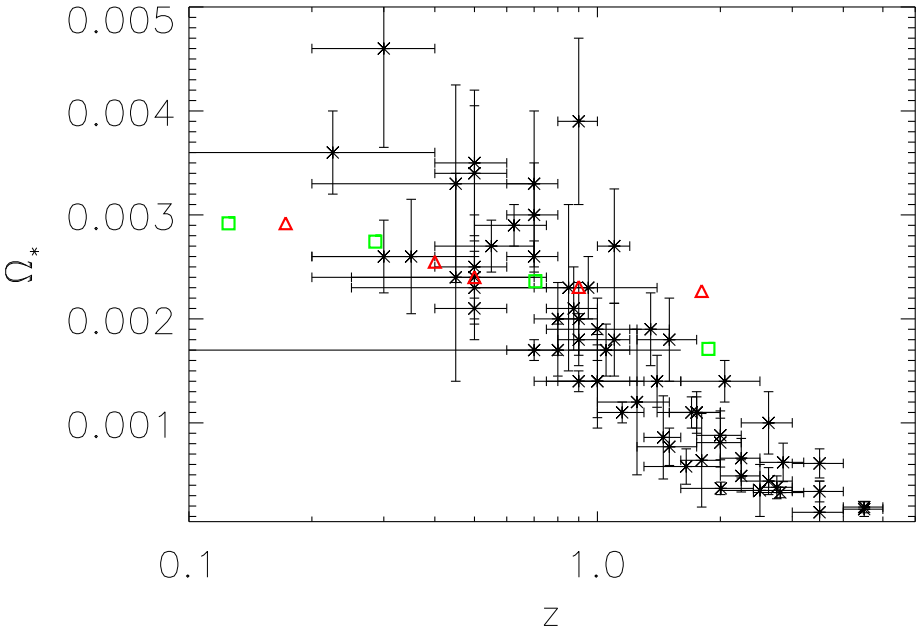

Figure 5.7: Red points: $\Omega_*(z)$ from the VESPA analysis. Green points: $\Omega_*(z)$ from the MOPED analysis (Panter et al. 2007). Black points: $\Omega_*(z)$ as compiled by Wilkins et al. (2008), matching the points of Figure 4.3.

Redshift	$\Omega_*(z)$ Salpeter	$\Omega_*(z)$ Chabrier
0.005-0.35	0.00292	0.00204
0.4	0.00240	0.00169
0.5	0.00231	0.00161
0.9	0.00227	0.00159
1.8	0.00234	0.00163

Table 5.8: $\Omega_*(z)$ as calculated by VESPA using equation (5.25).

majority of galaxies in the SDSS, a redshift of 1.8 falls in the last bin of star formation which, if fully resolved, spans 9 to 14 Gyr in the galaxy's rest frame. It is possible that this loss of resolution at high lookback is responsible for an overestimation of the mass at any given time within that bin. Even though we can constrain the total mass formed within any given bin, we have no information about *when* in the bin the mass formed. This problem is further aggravated in galaxies where the final bin has not been fully resolved.

The large width of the last bin is a reflection of the fact that old populations are more degenerate than younger populations, and it becomes increasingly hard to resolve their ages from an integrated galaxy spectrum. This sets a limit on how far back we can use VESPA's catalogue to infer the properties of galaxies at redshifts of 1.8 and higher. This limitation can be lifted by very high quality spectra, or higher redshift galaxy surveys.

5.3.4 IMF studies

Even though VESPA is only as good as the underlying models, the advantage is that VESPA can be used to compare different models with observations in an attempt to find out which models best fit the data. An example of this is the study of Wilkins, Hopkins, Trentham and Tojeiro (2008, in print), in which we attempted to conciliate the discrepancy observed between different measurements of the stellar mass density shown in Figure 4.3 and introduced in section 4.1.2.

The idea is to estimate the dependence of both the stellar mass density as calculated from integrating the cosmic star formation history over time, $\rho_{*,SFH}$, and the observed stellar mass density calculated by directly estimating the stellar mass of galaxies, $\rho_{*,obs}$, on the underlying IMF. The two methods are sensitive to changes in the IMF in different ways: instantaneous methods probe the presence of high mass stars, and stellar mass estimates are more sensitive to older, lower mass stars. We can therefore hypothesise that there is an IMF which brings the two measurements into agreement.

To test this we define a base IMF as a three-part power law:

$$\xi(M) \propto \begin{cases} M^{\alpha_0} & \text{if } 0.01 < M/M_\odot < 0.08 \\ M^{\alpha_1} & \text{if } 0.08 < M/M_\odot < 0.5 \\ M^{\alpha_2} & \text{if } 0.5 < M/M_\odot < 150 \end{cases} \tag{5.26}$$

with $\alpha_0 = 0.3$, $\alpha_1 = 1.3$ and $\alpha_2 = 2.35$. There is an infinite number of ways in which to parametrize an IMF, but a three-part power-law is in agreement with the local measurements (Kroupa 2007). In practice we want to estimate the dependence of $\rho_{*,SFH}$ and $\rho_{*,obs}$ on α_i and find the solution which minimises the difference between the two measurements. This relies on finding accurate calibration factors between masses or star formation rates across different values of α_i. Figure 5.8 shows how α_2 and α_1 affect the estimation of stellar masses in the two cases.

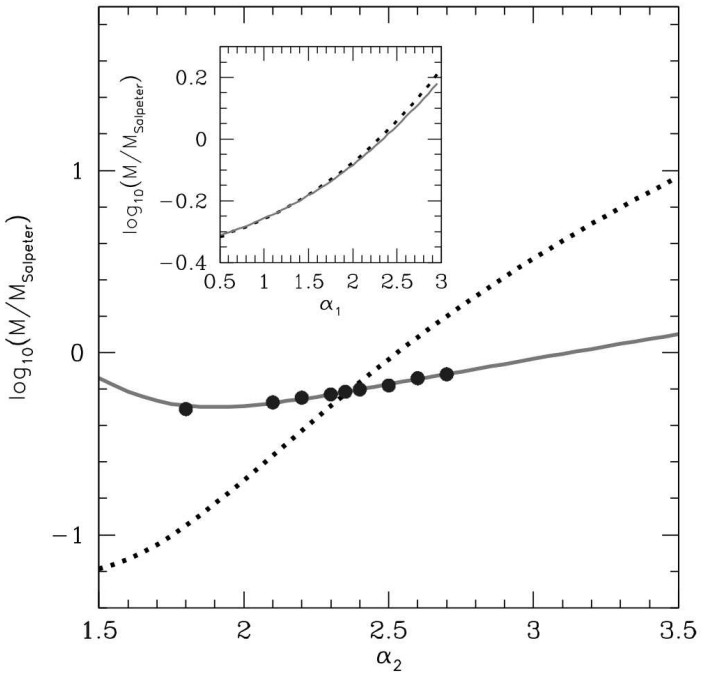

Figure 5.8: From Wilkins et al. (2008, in print). The main figure shows the dependence of $\rho_{*,SFH}$ (dotted line) and $\rho_{*,obs}$ (solid line and filled dots) on the high mass slope, α_2. The filled dots were obtained with VESPA - see main text for detail. The inset shows the dependence on the low mass slope, α_1.

The effect on $\rho_{*,SFH}$ is estimated by considering how a change in IMF changes the UV luminosity calibration normally used to infer an instantaneous star formation rate. This was done by considering PEGASE.2 (Fioc and Rocca-Volmerange 1999) stellar

population models with varying underlying IMFs, and the resulting changes in UV luminosity - the result can be seen in the dotted lines. The effects on other indicators used to estimate the instantaneous SFR, such as H_α or the rate of core-collapse SN have also been calculated. The dependency of stellar masses and $\rho_{*,SFH}$ on the IMF changes somewhat, but the general behaviour is the same.

The dependency of $\rho_{*,obs}$ on the IMF is harder to estimate given the widely different methods used to estimate stellar masses. One can start from a luminosity function, which combined with a mass-to-light ratio provides an estimate of the stellar mass density. This normally relies on near-IR colours, which are more representative of the underlying population as a whole (unlike UV light, which is dominated by emission from young stars). In this case, we can estimate how stellar mass estimates change with IMF by considering the changes in the near-IR light of stellar populations with a star formation history assumed to be that of the cosmic star formation history, as calculated by PEGASE.2 models - this can be seen in the solid lines. Another widely used method to estimate stellar masses is to use either broadband colours, spectral features or the full spectral range to estimate the star formation history given a set of models. To estimate the dependence of $\rho_{*,obs}$ on the IMF for masses estimated in this way, we have adapted VESPA to work with PEGASE.2 models, and have analysed a fixed sub-sample of 2000 galaxies with a range of IMFs. The resulting correction on the recovered masses gives the filled black circles in Figure 5.8.

We see that the high-mass slope affects the two measurements in significantly different ways, whereas the dependence on the low mass slope is very similar for the two cases. A changing high-mass slope affects $\rho_{*,SFH}$ more pronouncedly as expected from its dependence on light from young stars. From the figure alone we see that a flatter slope for high masses might bring the two local measurements in agreement. Using the star formation history of Hopkins and Beacom (2006) and the set of $\rho_{*,obs}$ compiled by Wilkins et al. (2008) we can apply the corrections derived above to obtain Figure 5.9.

Varying the value of α_2 only (given that this is the only slope to significantly affect the difference between the two measurements) and performing a χ^2 minimisation gives the area between the vertical lines in Figure 5.9. This corresponds to a best-fit value of $\alpha_2 = 2.15$. Even though this brings the local stellar mass density estimates into agreement, it does not solve the discrepancy at higher redshifts. We can repeat the analysis using all the available points in redshift, and ask for the value of α_2 which

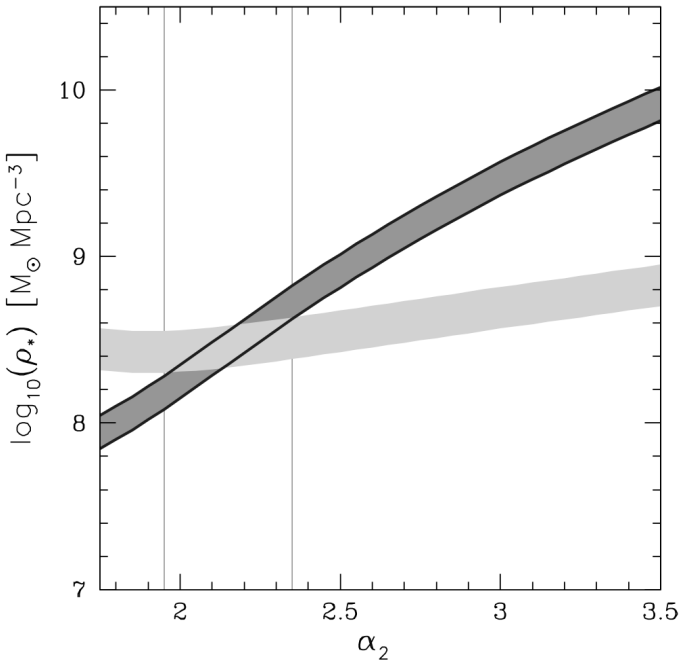

Figure 5.9: From Wilkins et al. (2008, in print). The predicted local stellar mass density as derived from integrating the cosmic star formation history (dark shaded area) and from direct stellar mass measurements (light shaded area) as a function of the IMF high-mass slope, α_2. The vertical lines give the range in α_2 which brings the two measurements into statistical agreement: $1.95 < \alpha_2 < 2.35$.

best fits the data. The answer is a shallower slope, $1.85 < \alpha_2 < 2.15$ with a best fit value $\alpha_2 = 2.0$. Whereas this is formally the model (within our parametrization) which best describes the data, it gives a formally poor fit, in particular at low redshift where the data is more trustworthy. This might be because the errors at high redshift are underestimated, or because there is not an IMF which follows our model which fits the data at all redshifts in a satisfactory way, or a combination of both.

We conclude from this study that an IMF as described by equation (5.26), with a modified high-mass slope $\alpha_2 = 2.15$ is able to reconcile measurements of $\rho_{*,SFH}$ and $\rho_{*,obs}$ at $z < 0.5$. However, a similar parametrization fails to find a model which simultaneously fits the data at low and high redshift. Systematic errors in either of the measurements could be the cause, or one might have to consider an IMF which has a dependence on some quantity which is in turn redshift dependent. Given the lack of full physical arguments to derive an IMF, constraints such as the ones presented in this

study are valuable observables which can drive theoretical models forward.

Chapter 6

Progenitors of supernovae type Ia

The observation of supernovae type Ia (SNIa) at a variety of redshifts has revealed the local Universe's expansion is accelerating (section 1.2.4 and references therein). Even though a range of cosmology experiments have confirmed that the best fit cosmological model is consistent with an accelerating Universe, SNIa experiments continue to provide the only *direct* piece of evidence for Dark Energy. Recent advancements on the measurement of the rate of SNIa has revealed that they are likely to have more than one type of progenitor star. This might have far-reaching consequences to both theoretical models of SNIa, as well as future Dark Energy experiments. In this Chapter we use the catalogue described in Chapter 5 to analyse the star formation histories of SNIa host galaxies, and extract information on their progenitor systems.

This chapter is organised as follows: in section 6.1 we discuss the current views and work on the type of SNIa progenitors and their evolution with redshift; in section 6.2 we present our method and in section 6.3 our results; in section 6.4 we discuss our findings and finally in section 6.5 we lay down the path for the future.

6.1 Introduction

SNIa are the product of a thermonuclear explosion of a white dwarf reaching its Chandrasekhar mass - the mass limit allowed by the supporting electron degeneracy pressure. This requires the white dwarf to accrete mass, and two models are proposed to explain

this: the *single-degenerate* model proposes that the mass is accreted from a companion main-sequence star; the *double-degenerate* model instead suggests the mass accretion comes from merging with another white-dwarf. Even though we currently lack a standard model for SNIa, the simple fact that they arise from the thermonuclear explosion of a star at a fixed mass suggests that they might have a tight distribution in brightness. In fact, this is not strictly true, and there is an intrinsic scatter in luminosity associated with these explosions. The success of SNIa as cosmological probes lies not in the fact that they are standard candles, but rather 'standardisable' candles.

Phillips (1993) showed that this intrinsic scatter on the peak luminosity of SNIa could be as high as 0.8 mag in the B-band, but more importantly they showed that there is a tight correlation between the peak luminosity and the decay-time of the light curve. This, and similar relations which correlate peak brightness with the width of the light-curve (the stretch factor), allow for the peak-magnitudes of a sample of SNIa to be empirically calibrated, and luminosity distances to be estimated with a small enough scatter as to allow us to differentiate between cosmological models. However, as illustrated by the panel (b) of Figure 1.1, a scatter remains in the Hubble diagram. If this scatter is purely statistical, then larger numbers of SNIa observations will always allow for a better constraint of cosmological parameters. However, as statistical errors are reduced in this way, we are likely to hit the wall of systematic errors. A possibility is that our model for the correction of peak-luminosities is too simple - in fact, it almost certainly is - and that the peak luminosity of any given SNIa depends on more than its decay time or light-curve shape. If it depends on something which is evolving with redshift, then we are introducing a systematic bias in the estimation of cosmological parameters.

What does this mean for cosmology? Future Dark Energy surveys have high hopes, which must be matched by high quality data. To keep systematic errors subdominant relative to statistical errors in a mission like the Supernova Acceleration Probe (SNAP), the scatter in magnitudes after correction should not be larger than 0.02 mag. This in itself demands a better understanding of SNIa properties (Ellis et al. 2008). Furthermore, if the properties of the supernovae change in a way as to introduce biased residuals which evolve with redshift, then this can introduce a bias in the estimation of the equation of state for Dark Energy (Howell et al. 2007; Sarkar et al. 2008).

The motivation for this project is two-fold. Firstly, given the information in the SFH of a galaxy, one can hope to learn about the type of progenitors of SNIa and put some constraints on their properties, such as mass, age and metallicity. This provides insight which can feed into and motivate SNIa stellar models. Secondly, VESPA has the potential to provide information on the most likely type of progenitor for a sample of SNIa. If the properties of progenitors are correlated with observables such as stretch or peak luminosity, which in turn might be correlated with residuals in the Hubble diagram, VESPA can contribute by identifying which stellar properties correlate with residuals, with a view to reducing them.

6.1.1 SNIa progenitors

The hypothesis that SNIa have more than one type of progenitor star is well motivated - brightest events (higher stretch) mostly happen in star-forming galaxies, whilst dimmer explosions (lower stretch) are associated with red, passive galaxies (Hamuy et al. 1996; Howell 2001; van den Bergh et al. 2005); the observed SNIa rate in blue galaxies is higher than that in red galaxies (Mannucci et al. 2005), and higher SNIa rates are also associated with radio-loud galaxies, when compared to a sample of radio-quiet galaxies (Della Valle and Panagia 2003).

Mannucci et al. (2006) explored this possibility. The delay-time distribution (DTD - $D(t)$) describes the range of delay-times between when a progenitor system is born and the creation of a SNIa explosion, for a given population of progenitors. The DTD therefore provides the connection between the observed SNR and the star formation history

$$SNR(t) = \int_0^t D(t - t')\psi(t')dt' \tag{6.1}$$

where $\phi(t')$ is the star formation rate at time t' and $D(t - t')$ can be read as the probability that a star of age t' will produce a SNIa at time t.

Mannucci et al. (2006) explicitly explored different shapes and types of DTDs, so as to better explain the observables mentioned in the previous paragraph which are inconsistent with a single-peaked DTD. They found that a two-peaked DTD is the best fit for their data, suggesting the presence of two distinct populations of progenitors.

The A+B model

A simpler and now popular model for these two populations is the so called A+B model, in which the observed SNR is described by two parameters - one proportional to the total stellar mass of the galaxy, and one proportional to the instantaneous star formation rate:

$$SNR(t) = AM_*(t) + B\dot{M}_*(t) \tag{6.2}$$

where $M_*(t)$ is the mass in stars and $\dot{M}_*(t)$ the star formation rate at time t. This formally corresponds to a DTD which is a δ-function at time t and associated with young stars (the *prompt* component), and a DTD which is constant in time and associated with older stars (the *delayed* component). Table 6.1 summarises some of the A and B rates measured in the literature. There is significant scatter across the different measurements for both rates. This is largely due to differences in the SNIa samples, methods for deriving A and B, definition for stellar masses and proxies for star formation rates.

Publication	A [SN $\mathrm{yr}^{-1} M_\odot^{-1}$]	B [SN $(M_\odot\ \mathrm{yr}^{-1})^{-1}\ \mathrm{yr}^{-1}$]
Neill et al. (2006)	$1.2 \pm 1.0 \times 10^{-14}$	$8.0 \pm 2.6 \times 10^{-4}$
Sullivan et al. (2006)	$5.3 \pm 1.1 \times 10^{-14}$	$3.9 \pm 0.7 \times 10^{-4}$
Scannapieco and Bildsten (2005)	$4.4^{+1.6}_{-1.4} \times 10^{-14}$	$2.6 \pm 1.1 \times 10^{-4}$
Scannapieco and Bildsten (2005)	$4.4^{+1.6}_{-1.4} \times 10^{-14}$	$1.2^{+0.7}_{-0.6} \times 10^{-4}$
Dilday et al. (2008)	$2.8 \pm 1.2 \times 10^{-14}$	$9.3^{+3.4}_{-3.1} \times 10^{-4}$

Table 6.1: Summary of A and B rates calculated in the literature. The two B rates from Scannapieco and Bildsten (2005) are calculated using two different estimators for the star formation rate.

Figure 4.1 shows that the cosmic star formation rate has evolved strongly with redshift. This suggests that if equation (6.2) is a good model for the SNIa rate, then we should expect the weight of each of the routes to SNIa to change significantly with redshift. This is well illustrated in Figure 6.1, in which the contribution of the two components is estimated as a function of redshift. As expected, the prompt component is dominant at higher redshifts when the SFR was higher and declines in more recent times, whereas the delayed component shows the opposite behaviour. If the two populations are intrinsically different, one can ask the question of whether a single empirical calibration relation applies. If it does not, then calibrating peak luminosities of SNIa at higher redshifts using empirical relations estimated at low redshift (or

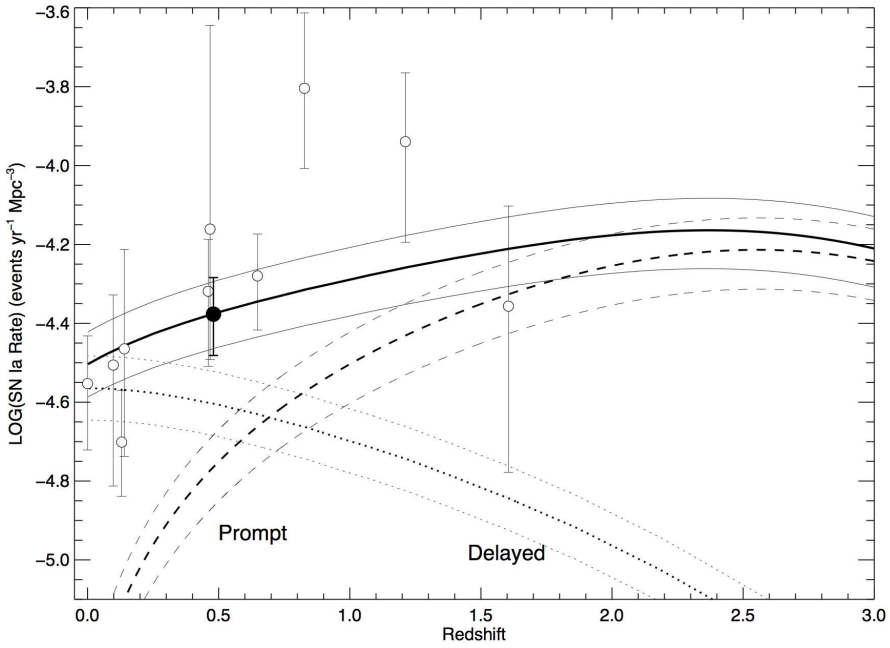

Figure 6.1: From Sullivan et al. (2006): Prediction of SNR as a function of redshift, assuming equation (6.2). A and B were estimated using the SNLS sample at $0.2 < z < 0.75$ and assumed not to evolve with redshift. The star formation rate history of Hopkins and Beacom (2006) was assumed.

vice-versa) will introduce systematic errors which have a strong dependence on redshift.

Beyond the A+B model

There have also been attempts to explicitly measure the DTD from the data. Pritchet et al. (2008) use the relationship between the SNR and SFR to find $D(t) \propto t^{-0.5\pm0.2}$. Using a theoretical argument to calculate the rate of white-dwarfs per unit stellar mass, they find that a single-degenerate model can yield a DTD of this form only if the fraction of white dwarfs which explodes as SNIa (constrained in that paper to be 0.01) is independent of mass. Given that this is unlikely to be the case, they conclude that there must be another route to SNIa. Totani et al. (2008) find $D(t) \propto t^{-1}$ in the range $0.1 < t < 10$ Gyr, by looking at the ages of SNIa hosts in a sample of passive red galaxies, but cannot put any constraints on the DTD of SNIa which occur in star-forming galaxies. The wide range of DTD models allowed by the data at this stage can be seen simply as an indication that the data is not of yet good enough to allow a clear

distinction.

6.1.2 SNIa evolution

If there are more than one progenitor to SNIa, or if SNIa properties evolve with redshift for one reason or another, then one can hope to find changes in the properties of SNIa at low and high redshifts. Hook et al. (2005), using 14 spectra of SNIa with $0.17 < z < 0.83$ from the Supernovae Cosmology Project found no evidence for evolution, but due to the small size of the sample the only conclusion is that there is a population at high redshift which has similar spectral properties to those at low redshift. Bronder et al. (2008), based on a larger sample of 87 SNIa from the SNLS sample ($\tilde{z} = 0.720$) which they compare to a local sample from the literature also find that the two samples have spectral properties which are consistent with a single population. Foley et al. (2007) construct composite spectra of low ($0.15 < z < 0.77$) and high ($z > 1$) SNIa, which they compare. Even though they found the two samples to be remarkably similar, they found some minor deviations. However, the interpretation of these deviations is not straightforward, mainly due to a lack of theoretical support. Even though they found an excess in the UV (a good probe of the progenitor's metallicity) in the high-redshift sample, different models offer opposite interpretations for the significance of this excess.

Ellis et al. (2008) look at the spectral properties in the UV of a sample of 36 SNIa at an intermediate redshift ($\bar{z} = 0.5$), which they compare to a local sample. Again, no evidence for evolution was found. Interestingly, they find a large scatter in the UV ($\lambda < 3000 - 4000$ Å) spectrum from one SNIa to another which is inconsistent with metallicity changes offered by current models, or changes associated with dust extinction. This scatter remains even after colour corrections, which normalise the colours at other wavelengths. The lack of theoretical understanding of this scatter is worrying. At high-redshift surveys, cross-colour k-corrections are necessary in order to estimate rest-frame light curves, and this requires a template that is reliable in the UV. The observed scatter in the UV colours introduces on itself a statistical uncertainty of 0.05 to 0.1 magnitudes. This need not be a concern - if these residuals are randomly distributed around the template then one needs only to increase the size of the SNIa sample. However, they also find that the average UV spectrum of low and high stretch objects is different. This introduces the risk of using the wrong template to calculate cross-colour corrections. Given the correlation between high-stretch with bright explosions in star forming galaxies, and low stretch explosions with dimmer SNIa in passive

galaxies, the A+B model predicts an evolution of the average stretch with redshift. This was observed by Howell et al. (2007), who found an increase in the average stretch with redshift that is consistent with the A+B model.

6.2 VESPA and SNIa

Given the star formation history of a sample of SNIa host galaxies and the observed SNR, one can hope to recover the DTD of the progenitor population(s) - the supernovae rate at any time t is a convolution of the star formation history as a function of time, with the DTD (equation 6.1).

Given the lack of direct observations of SNIa progenitors, the DTD distribution provides an important link between theoretical models and observations. This is to say, the information about the progenitors of SNIa of a given sample is encoded in its star formation history.

6.2.1 Sample selection

We gathered a sample of around 1300 confirmed SNIa from IAU circulars[1], the CfA supernovae list[2] and the SDSS-SN public list of supernovae (Frieman et al. 2008), which we cross-matched with the SDSS DR5 (Adelman-McCarthy et al. 2007) spectroscopic sample. This yielded a list of 256 galaxies which are known to be hosts to 257 SNIa (one galaxy had two SNIa). Even though this number is not enough to attempt a full deconvolution of equation (6.1), we can use the star formation history of the hosts to test the existence of a short-lived progenitor to SNIa and to put a constraint on its age, which in turn allows us to put a constraint on its mass.

As a control sample we use a randomly selected set of SDSS galaxies from DR5, weighted as to provide the same redshift distribution as the hosts sample.

6.2.2 Method

We follow a parametrization similar to equation (6.2):

$$SNR = \alpha M_* + \beta M_{recent} \tag{6.3}$$

[1] http://www.cfa.harvard.edu/iau/cbat.html,

[2] http://www.cfa.harvard.edu/iau/lists/Supernovae.html,

where SNR is the observed SNIa rate, M_* is the current stellar mass in a galaxy and M_{recent} is the current stellar mass in stars formed at $t = \Delta t_{recent}$. We can re-write equation (6.3) as

$$\frac{SNR}{M_*} = \alpha \left(1 + \frac{\beta}{\alpha}\frac{M_{recent}}{M_*}\right). \tag{6.4}$$

Given the way in which we collect our sample, we have no way to estimate absolute efficiencies of the supernovae rates. However, we can estimate the quantity $\frac{SNR}{M_*}$ as a function of $\frac{M_{recent}}{M_*}$, up to an unknown constant of proportionality and from there infer the value of the quantity β/α.

The observed SNR, per unit of stellar mass, is simply

$$\frac{SNR}{M_*} = \epsilon\frac{N_{SN}}{M_{total}} \tag{6.5}$$

where ϵ represents the SNIa detection efficiency for the sample, which is unknown to us. For a given value of Δt_{recent} we do the following:

1. We divide the hosts and control sample into sub-samples according to their value of recent star formation, $\frac{M_{recent}}{M_*}$, as derived by VESPA.

2. We estimate $\frac{SNR}{M_*} = \epsilon\frac{N_{SN}}{M_{total}}$ for the galaxies in each sub-sample. M_{total} is estimated from the much larger control sample.

3. We use the relation (6.4) to find the value of β/α which best fits the data.

In doing so we are making two key assumptions: a) we are assuming that ϵ, the detection efficiency, is independent of host type, and b) we assume that our sample is an un-biased sample of SDSS galaxies. We discuss these two issues in section 6.4.

VESPA provides both the stellar masses and the star formation fractions necessary for this project. Given that the hosts' spectra varied in quality, we choose to use a one-parameter dust model to analyse both the hosts and the control sample, as the data often does not call for a more sophisticated model.

We choose values of Δt_{recent} which match VESPA's bin boundaries. This is the natural choice, even though it limits our resolution somewhat. For a given value of Δt_{recent} we can calculate $\frac{M_{recent}}{M_*}$ in the corresponding bins for the hosts and control sample. Together with equation (6.5), we can then constrain the relation given by

equation (6.4) and get a value for β/α. Given the hypothesis that there is a component of the SNR which is proportional to the amount of young stars, we expect β/α to be highest for Δt_{recent} of the order of the age of the progenitors. On the other hand, a value of $\beta/\alpha = 0$ corresponds to zero contribution to the SNR from stars formed in Δt_{recent}.

6.3 Results

Each galaxy has its star formation history cast into a different configuration of bins. As mentioned in section 4.2.3, each VESPA iteration begins with 3 low-resolution bins, the youngest of which spans the ages from $t = 0.02$ to $t = 0.74$Gyr. This means that this boundary is resolved for all galaxies. This is an advantage because for unresolved bins we cannot pin down when in the bin the recovered mass formed. For $\Delta t_{recent} = 74$ Myr we find $\beta/\alpha = 465 \pm 83$, meaning that stars of this age have a significant contribution to the observed SNR. This has important consequences in itself, given that the masses associated with such short-lived progenitors are very high ($> 5.5 M_{\odot}$).

To investigate how older stars contribute to the SNR we include older bins in our analysis. We find that β/α is significant up to 180 Myr ($\beta/\alpha = 454 \pm 78$) after which the correlation decreases. Stars aged between 180 and 660 Myr contribute very little to β/α, as do even older stars. Table 6.2 and Figure 6.2 summarise our results. Given the decrease in contribution from stars older than 180 Myr to β/α we put a constraint of 180 Myr on the age of progenitor systems which are associated with the prompt route.

There is a subtlety here with respect to the fact that the boundary at 180 Myr is not one which is necessarily resolved for all objects. For the host sample however, we find that the vast majority of hosts do indeed resolve this boundary. We also do not find any significant difference in our results if we exclude the hosts which do not.

6.4 Interpretation

We find a significant correlation between a SNIa and stars formed in the last 180 Myr, which we identify as the progenitors of SNIa formed by the prompt route. This correlation decreases significantly in older bins, and we therefore, for the first time, put

Δt_{recent} (Myr)	β/α	Error on β/α	Significance
0 - 180	454	78	5.8σ
180 - 660	56	16	3.4σ
660 - 2440	18.4	3.5	5.2σ
2440 - 13700	-3	1.0	-

Table 6.2: Best-fit values of β/α for different aged populations and the respective significances. We see that there is a high correlation between SNIa and stars younger than 180 Myr, which decreases significantly when we move to older stars. We interpret this as the prompt route to SNIa, and postulate that the progenitors associated with this route are younger than 180 Myr.

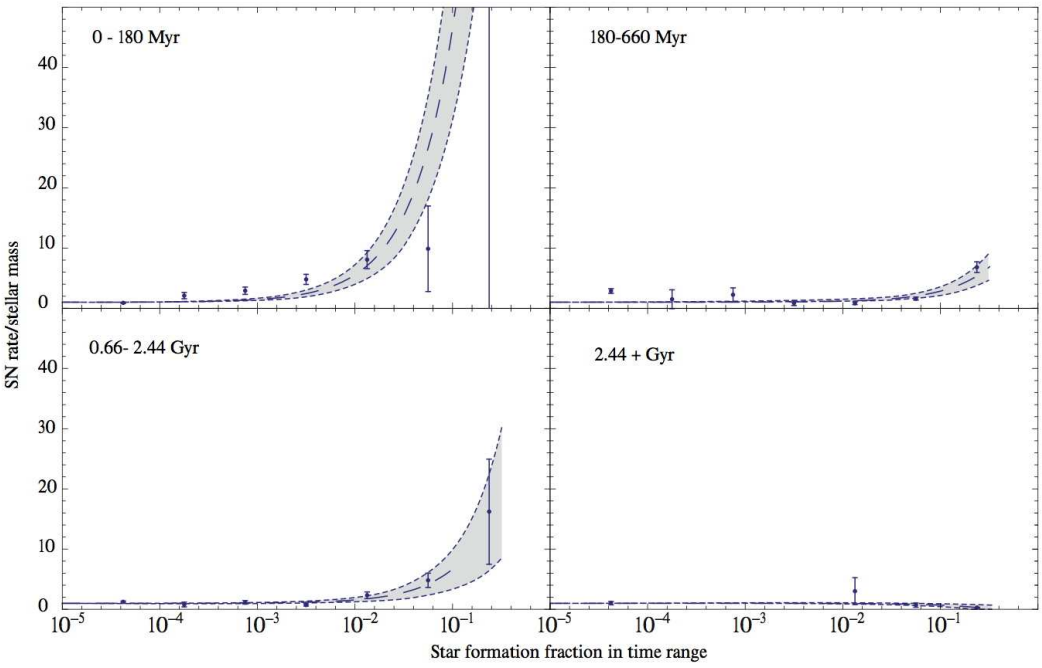

Figure 6.2: SNIa rate per unit stellar mass, unnormalised, as a function of fraction of stellar mass formed in $\Delta t = [0-180], [180-660], [660-2444]$ and $[2444, 13700]$Myr. We see that there is a high correlation between SNIa and stars younger than 180 Myr, which decreases significantly when we move to older stars. We interpret this as the prompt route to SNIa, and postulate that the progenitors associated with this route are younger than 180 Myr.

an upper limit on the age of the progenitor stars associated with the prompt route.

6.4.1 Sample selection effects

Our result would be mimicked if our hosts sample was biased towards high recent star formation for reasons *other* that SNIa related physics. This might happen if for example the SNIa came from targeted surveys which preferentially pick out events in blue galaxies. The only targeted survey which contributes towards our sample is the Lick Observatory Supernova Search (LOSS, Richmond et al. 1998). Removing the SNIa from this survey from our analysis changes our results without significance. We also find no significant difference in the distribution of dust values between the host and control sample, meaning the hosts sample is not biased towards low or high dust values.

A similar issue is the question of whether the efficiency of detection ϵ can introduce such a bias. We might expect the detection efficiency to depend on the host type: fainter (low stretch) SNIa, typically from massive red galaxies are harder to detect, especially at high redshifts. If this was a strong effect we should expect to see our hosts sample deficient in high-mass galaxies. In fact we see the opposite behaviour - the stellar mass distribution of the hosts sample is skewed to high masses. This is not unexpected: the SNR is proportional to stellar mass, which partially confuses the picture and does not allow us to conclude that this sort of variation in efficiency is irrelevant for our analysis. A quick sanity check is to re-do the analysis using only hosts with $z < 0.1$, where we might expect this effect to be least pronounced. This is certainly important for the SDSS SNIa, for which the efficiency is less than one for $z > 0.1$ and where they miss a higher proportion of the faint hosts (Dilday et al. 2008). Doing this analysis reduces the sample to 190 hosts, but the result remains consistent with our initial measurement: $\beta/\alpha = 514 \pm 104$.

6.4.2 Comparing with previous results

Given that we cannot calculate absolute α and β rates, we cannot directly compare our results to others in the literature. However, if we transform the B rate in the literature - usually in SN $M_{\odot}^{-1}$ - to SN $M_{\odot}^{-1}\mathrm{yr}^{-1}$ by dividing by 180Myr, we can compare published values of B/A with our β/α. Clearly, given that the instantaneous SFRs used to calculate B were averaged over different time scales means that even this comparison is hard to interpret. Using the values from Neill et al. (2006) gives B/A

≈ 300. Using the same method, values from Sullivan et al. (2006) yield B/A ≈ 40 and those from Scannapieco and Bildsten (2005) give 300 and 150.

6.4.3 On the nature of the progenitors

Such a short time delay strongly constrains the nature of possible progenitors. They must be stars that evolve fast enough, i.e. with a mass above $\sim 3.5 M_{\odot}$, but must be below the super-AGB mass threshold ($\sim 8 M_{\odot}$) above which one gets electron-capture supernovae (Poelarends et al. 2008). Pinsonneault and Stanek (2006) have also suggested that a significant fraction of binaries are twins (i.e. pairs of stars with essentially identical masses), and that such twin binaries could produce a short (< 0.1 Gyr) path to SNIa. Considering common envelope evolution phenomena, Pinsonneault and Stanek (2006) argue that such twin systems could yield double degenerate SNIa in a way that would be both fast and efficient.

Are there enough high-mass progenitors to account for the observed SNIa rate? Only a fraction of these stars f_{β} will actually explode as a SNIa. We take into account five factors: the fraction of stars in binaries (f_a), the fraction of the binaries both of whose components lie in the range 3.5 to 8 $M_{\odot}$ (f_b), the fraction of stars at a suitable separation for mass transfer (f_c), the fact that every binary yields a single explosion, (f_d), and an overall efficiency (η_{β}, as not all possible progenitors may explode). Maoz (2008) has estimated the first four factors, and finds $f_a \in [2/3, 1], f_b \in [1/6, 1/3], f_c \in [1/4, 1/2]$ and $f_d = 1/2$. Crudely multiplying these factors together gives the fraction of objects in the appropriate mass range that explode as prompt SNIa: $f_{\beta} \in [0.014, 0.083]\eta_{\beta}$.

The fraction of stars which will explode as a SNIa progenitor, is given by

$$f_{\beta} = \frac{N_{\beta}}{N_{3.5-8}} \tag{6.6}$$

where N_{β} is the total number of SNIa from the fast route, and $N_{3.5-8}$ is the total number of stars in the correct mass range. Our result only allows us to estimate the ratio of the two components. We can however use published values of A in order to infer an absolute value of β from our ratio - we call this value B'. Using the A value from Neill et al. (2006) gives $B' = (5.4 \pm 4.2) \times 10^{-12} SN\,\mathrm{yr}^{-1}\,\mathrm{M}_{\odot}^{-1}$. We can now estimate $N_{\beta} = B' \times 180\,\mathrm{Myr} \times M_{180}$ and $N_{3.5-8} = 0.0157 \times M_{180}$, assuming a

Salpeter IMF of the form $N(M) \propto M^{-2.35}$. This gives $f_\beta = 0.063 \pm 0.048$, within the values given by Maoz (2008). Using the A value from Sullivan et al. (2006) gives a value of $B' = (2.41 \pm 0.65) \times 10^{-11} SN\,\mathrm{yr}^{-1}\,\mathrm{M}_\odot^{-1}$ and $f_\beta = 0.28 \pm 0.07$. Similarly, A as published by Scannapieco and Bildsten (2005) results in roughly $B' = (1.99 \pm 0.69) \times 10^{-11} SN\,\mathrm{yr}^{-1}\,\mathrm{M}_\odot^{-1}$ and $f_\beta = 0.23 \pm 0.08$. With the slow A rate of Neill et al. (2006) there is complete consistency with the theoretical expectations of SNIa rates from Maoz (2008). However, with the higher A rates of Sullivan et al. (2006) or Scannapieco and Bildsten (2005), there is some tension with our results. It is hard to know exactly how much meaning to extract from this: firstly we are relying on external measurements of the A rate, which show significant scatter themselves. Secondly an excess of SNIa explosions with respect to the predicted number of progenitors is observed in a large number of SNIa studies (Maoz 2008). This tension can be alleviated in a number of ways - some of them discussed in the above paper - but generally indicates that the efficiency of the mechanism which produces SNIa explosions must be very high.

6.5 Future prospects

With the current size of the hosts sample we were already able to put the first tight constraint on the age of the progenitor stars associated with the prompt route.

We are currently analysing a larger sample of SNIa hosts, which will allow us to extract more information on the DTD. A tight measurement of the DTD would in itself be an extremely valuable observational constraint on theoretical models of SNIa.

With the public release of the SDSS SNIa sample, we will investigate how residual scatter around the Hubble diagram correlates with physical parameters such as progenitor age and metallicity. VESPA's contribution here is essential, as we need robust star formation and metallicity histories of the hosts. Ultimately we hope to develop a method which contributes towards a better understanding of SNIa explosions, light-curve calibrations, and the nature of Dark Energy.

Chapter 7

Summary and future work

This thesis has presented work done on a variety of scales and epochs and explores several aspects related to the evolution of the Universe. Even though we have presented a summary and our conclusions at the end of each relevant section or Chapter, for the sake of completeness and clarity we take the opportunity to expand and collect all of our main conclusions and discussion in this final Chapter.

7.1 Non-Gaussianity studies of the CMB

We first looked at non-Gaussianity studies of the cosmic microwave background. Inflation predicts Gaussianity in the temperature fluctuations, which we analysed using the peak-peak auto- and cross-correlation of temperature peaks. We studied the first- and fifth-year data releases, and looked for signs of non-Gaussianity.

In the first year analysis, using the auto-correlation of peaks, we detected a signal that was associated with cold spots in the southern hemisphere. Removing the regions of sky within 30 degrees of the galactic plane, or scales with $\ell < 40$, removed the non-Gaussian signal. An analysis of individual frequencies showed that the Q-band was the most affected by this signal, but not the only one. We concluded that the signal we see is most likely due to the presence of unsubtracted or oversubtracted foregrounds.

We revisited the problem with the fifth year data, and introduced a number of improvements over our first year analysis: in addition to the auto-correlation we also calculated the cross-correlation of peaks across different frequencies and we introduced computational improvements which allowed us to analyse more peaks, as well as reach better angular resolution. Once again we found a signal that was associated with cold spots in the southern hemisphere, and the Galactic plane. Our interpretation is that this is the same signal we saw in the first year analysis, which we again attribute to a non-cosmological origin. In addition to this signal, however, we are puzzled by a detection which appears only in the northern hemisphere, outwith 30 degrees of the galactic plane, in the V and W band. This signal is robust to the choice of binning, matrix-invertion methods and number of simulated maps, and it does not seem to be associated with any particular angular scale on the sky. Further investigation is required to pin down its origin. The next obvious steps are to increase our constraints in real space, and also in harmonic space, which we plan to do. We also observed some anomalies associated with hot spots, mostly confined to the northern hemisphere.

We explored the possibility that the local integrated Sachs-Wolfe (ISW) effect is the reason for these detections. We removed a physically-motivated ISW map, reconstructed from matter-density estimated to $z = 0.3$, from the WMAP data but did not find it to consistently resolve any of our detections. Our conclusion from this analysis is that the ISW is not the reason behind the non-Gaussianity we see.

Finally we note that the first and fifth year analyses were done using a different set of peaks, each selected according to a temperature threshold which was different in the two analyses. A lower threshold does in principle make us more susceptible to residual point-sources contamination, given that strong point sources are more likely to have been detected and masked. This might be the reason behind our hot spots detections in the V-band, although further investigation is required.

Our conclusion from the first and fifth year analyses, principally taking into account the frequency-dependence of our detections, is that there is substantial evidence for a temperature component in the WMAP data which is not primordial. Based on our analysis we suggest this component is strongest in the southern hemisphere, in the regions close to the Galactic plane.

Using non-Gaussian simulated maps, we investigated how sensitive peak statistics are to an f_{NL} type of non-Gaussianity. This is in principle a much more rewarding and exciting way to probe non-Gaussianity, as there is a clearer relationship between a signal and a physical cause. We note, however, that a detection of a positive f_{NL} could be misinterpreted if the temperature fluctuations are tainted with non-cosmological anomalies. The analysis given above suggests that we must still be careful when drawing conclusions about the physical origin of f_{NL} detections in the current WMAP data.

Using a low threshold for peak-selection (corresponding to a large number of points being selected) we are drawn to conclude that peak statistics are not a competitive estimator. However, we have argued that the threshold selection might be an important factor when it comes to sensitivity to f_{NL}. We have in fact already explored a range of temperature threshold values, as well as high-pass filters to help with the problem of cosmic variance. However, after roughly two months of computational time, the results were lost due to serious hardware failure. The simulations will be re-done once new harware comes online, which will allow us to investigate what the optimal threshold and filter are and finally apply our method to WMAP data, as well as Planck data in the future.

7.2 VESPA

A significant amount of time during this PhD was devoted to the development and testing of VESPA - a new algorithm for recovering star formation and metallicity histories from galactic spectra. We were motivated by current limitations of other methods, which use a single parametrization for every analysed spectrum, independent of its quality. Whereas this is a perfectly valid approach when analysing large ensembles of galaxies, we wanted to have a method which would reflect the quality of the spectrum in the solutions it provides, on a galaxy-by-galaxy basis.

After significant tests with synthetic and real data, we found that VESPA is very successful at doing this. Particularly, we note that the self-regularization that we impose does a much better job at finding a solution which is closer to the true solution than one without such a regularization imposed. This will always happen at the expense of goodness of fit in data space, but that is the nature of the game - we sacrifice

precision for the sake of accuracy.

Our analysis of SDSS spectra showed that we can, on average, recover between two and five populations, and very rarely did we get more than six. We should note however, that this number is attached to the binning we chose. This binning already reflects the fact that older populations are more degenerate, so our rule of thumb of 2-5 populations would not apply if one was interested exclusively in a population older than, for example, 5Gyr.

One of the main contributions of this project to the community is the publication of a catalogue with robust star formation and metallicity histories, stellar masses, dust values and other derived quantities. The catalogue is easily cross-matched with SDSS's database, opening up the opportunities for almost instant science.

As any other method or algorithm, VESPA is in a state of constant improvement. Future work in this area will include:

- adding IR and UV photometry to SDSS spectra and investigating how the recovered quantities are affected,
- exploring how much information one can get using photometry alone, and
- exploring how much information is lost by removing the large-scale information in the spectrum, which is often riddled with systematics.

Finally we should remind ourselves that VESPA's analysis is ultimately model dependent, for which the importance of up-to-date models and an understanding of their successes and shortcomings cannot be overstated.

7.3 Mass assembly in galaxies

With robust stellar masses and star formation histories for DR5's 465,000 galaxies, we explored the galaxy stellar mass function and stellar mass density in SDSS. As others had done before, we found a galaxy stellar mass function (GSMF) that is well fit by a Schechter function and found no evidence for evolution in the GSMF between the redshifts of 0.005 and 0.35 for the high mass range.

We took this analysis one step further and attempted to put a constraint on the role of mergers to the mass assembly of present-day galaxies. The idea is on paper very simple. We aim to disentangle growth due to star formation from growth due to merging by evolving each galaxy back in time, and find how many of its stars had formed at any given redshift. Integrating this knowledge over all galaxies tells us how much stellar mass had formed by that redshift, per unit volume. What it does not tell us is over how many galaxies that stellar mass was distributed. We therefore need an external measurement, which provides the final piece of information. This is the GSMF observed directly at that redshift and it tells us what the mass distribution of galaxies was like at that time.

We found that the data are currently not good enough to conduct this measurement. We look forward to future redshift surveys, which will allow us to analyse galaxies at low and high redshifts in a consistent way. This should provide us with the opportunity for the first clean measurement of galaxy growth for reasons other than star formation, as a function of mass.

The opportunity remains to do a model-dependent version of this analysis, i.e. given a merging rate per galaxy stellar mass, is the observed GSMF compatible with the inferred GSMF from VESPA? Our feeling is that the data are presently not good enough to differentiate between theoretically interesting models, but that will change once new surveys become available.

We also looked at the evolution of the stellar mass density, which does not require any knowledge beyond the information that we have. Our results are consistent with previous estimates all the way up to redshift of 1.8. Here VESPA predicts a higher stellar mass density than those published in the literature, which we interpret as a limitation of the modelling. We simply cannot reliably distinguish between periods of star formation of ages higher than 9 Gyrs. In many cases, and depending on the quality of each spectrum, VESPA will say that this range is even more limiting and we cannot robustly pin down events of old star formation; the fossil record can only go so far. High-redshift spectroscopic surveys or better quality data are the key here, and we expect these degeneracies to be lifted with future datasets.

The evolution of the stellar mass density is not only an important observational

constraint for galaxy formation models. Recently it has been thrown into the spotlight due to unexpected discrepancies between different measurements. The reason for this discrepancy could be as simple as systematic effects, but it might also be telling us about the shape and potential evolution of the IMF. We used VESPA to estimate how stellar mass estimates are affected by different IMFs and investigate whether there is an universal IMF which resolves this discrepancy. The data tell us that there is not a universal power-law IMF which is able to do so. One exciting possibility is that the IMF has evolved with redshift, but one should consider this possibility with caution. The implications for galaxy formation would be substantial, and we must first make sure we understand all systematic effects associated with these measurements.

7.4 Progenitors to SNIa

One of the most recent additions to the collection of mysteries of the Universe is Dark Energy. It came as a true surprise, 10 years ago, that the Universe is now dominated by an energy component characterized by a negative pressure. Finding out more about its nature has become the primary goal of many scientists and scientific experiments.

SNIa experiments are particularly good probes of Dark Energy, but they rely on empirical calibrations of their light-curves. The suitability of these calibrations over a large range of redshifts has recently been questioned, and the possibility of a calibration law which evolves with redshift has been raised. This could lead to a higher uncertainty in the measurement of w, a bias, or both. The problem here is that we do not yet have a theoretical model for SNIa progenitors. Observational constraints on the nature of SNIa progenitors are badly needed both by those who study SNIa and those who use them as cosmological probes.

By studying the star formation history of SDSS galaxies which are known to be SNIa hosts, we were able to put the first tight constraint on the age of one type of progenitor, at 180 Myr.

This work has enormous potential for the future. We are already analysing a larger sample of SNIa hosts, which in itself will tell us more about the full range of progenitors to these important events. We also plan to use future SNIa datasets in order to look for a correlation between residuals on the Hubble diagram and physical properties of the

progenitors which can be assigned - at least in a statistical sense - by VESPA. Ideally one would like to do this one step up in the pipeline - i.e., to correlate residuals of the *light-curve* fit with progenitor properties, but this depends of the availability of suitable datasets.

7.5 Final remarks

It is nothing short of inspiring that we stand here today with a cosmological model that can take us from primordial perturbations in a scalar quantum field to a Universe populated with a wide range of galaxies. Whereas the big picture certainly seems to be in place, attention is now divided between fundamental questions such as the nature of Dark Energy, Inflation and structure formation, and the slightly messier non-linear physics such as star formation in galaxies, stellar mass assembly and all the feedback processes involved. This thesis is a small reminder that the two are intrinsically related.

As we approach the era of the great cosmology experiments, we are reminded of the fact that beating down statistical errors with larger surveys has its own natural barrier. We take the example of SNIa experiments, which now demand a true understanding of the physical processes behind these explosions and the host galaxies they inhabit if we are to interpret the data with confidence. The same is true in CMB analysis. Planck will be close to a perfect experiment, but non-Gaussianity studies - which will be yet again thrown into a new era with Planck - will always be plagued by the suspicion of a non-primordial temperature component in the data. In these two cases, and certainly in others too, the solution must be a combination of thorough and relevant observations with a better theoretical framework. As the physics gets more and more complicated, observations tend to, but do not always, lead the way.

This thesis makes intellectual and practical contributions in some of these areas, and paves the way for more and exciting future work. I hope the results and methods presented here will continue to contribute towards a better understanding of the formation, evolution and ultimate fate of our Universe.

Bibliography

Abazajian, K. et al.: 2005, AJ **129**, 1755

Abraham, R. G. et al.: 2007, *preprint (astro-ph/0701779)*

Adelman-McCarthy, J. K., Agüeros, M. A., Allam, S. S., Allende Prieto, C., Anderson, K. S. J., Anderson, S. F., Annis, J., Bahcall, N. A., Bailer-Jones, C. A. L., Baldry, I. K., Barentine, J. C., Bassett, B. A., Becker, A. C., Beers, T. C., Bell, E. F., Berlind, A. A., Bernardi, M., Blanton, M. R., Bochanski, J. J., Boroski, W. N., Brinchmann, J., Brinkmann, J., Brunner, R. J., Budavári, T., Carliles, S., Carr, M. A., Castander, F. J., Cinabro, D., Cool, R. J., Covey, K. R., Csabai, I., Cunha, C. E., Davenport, J. R. A., Dilday, B., Doi, M., Eisenstein, D. J., Evans, M. L., Fan, X., Finkbeiner, D. P., Friedman, S. D., Frieman, J. A., Fukugita, M., Gänsicke, B. T., Gates, E., Gillespie, B., Glazebrook, K., Gray, J., Grebel, E. K., Gunn, J. E., Gurbani, V. K., Hall, P. B., Harding, P., Harvanek, M., Hawley, S. L., Hayes, J., Heckman, T. M., Hendry, J. S., Hindsley, R. B., Hirata, C. M., Hogan, C. J., Hogg, D. W., Hyde, J. B., Ichikawa, S.-i., Ivezić, Ž., Jester, S., Johnson, J. A., Jorgensen, A. M., Jurić, M., Kent, S. M., Kessler, R., Kleinman, S. J., Knapp, G. R., Kron, R. G., Krzesinski, J., Kuropatkin, N., Lamb, D. Q., Lampeitl, H., Lebedeva, S., Lee, Y. S., Leger, R. F., Lépine, S., Lima, M., Lin, H., Long, D. C., Loomis, C. P., Loveday, J., Lupton, R. H., Malanushenko, O., Malanushenko, V., Mandelbaum, R., Margon, B., Marriner, J. P., Martínez-Delgado, D., Matsubara, T., McGehee, P. M., McKay, T. A., Meiksin, A., Morrison, H. L., Munn, J. A., Nakajima, R., Neilsen, Jr., E. H., Newberg, H. J., Nichol, R. C., Nicinski, T., Nieto-Santisteban, M., Nitta, A., Okamura, S., Owen, R., Oyaizu, H., Padmanabhan, N., Pan, K., Park, C., Peoples, J. J., Pier, J. R., Pope, A. C., Purger, N., Raddick, M. J., Re Fiorentin, P., Richards, G. T., Richmond, M. W., Riess, A. G., Rix, H.-W., Rockosi, C. M., Sako, M., Schlegel, D. J., Schneider, D. P., Schreiber, M. R., Schwope, A. D., Seljak, U., Sesar, B., Sheldon, E., Shimasaku, K., Sivarani, T., Smith, J. A., Snedden, S. A.,

Steinmetz, M., Strauss, M. A., SubbaRao, M., Suto, Y., Szalay, A. S., Szapudi, I., Szkody, P., Tegmark, M., Thakar, A. R., Tremonti, C. A., Tucker, D. L., Uomoto, A., Vanden Berk, D. E., Vandenberg, J., Vidrih, S., Vogeley, M. S., Voges, W., Vogt, N. P., Wadadekar, Y., Weinberg, D. H., West, A. A., White, S. D. M., Wilhite, B. C., Yanny, B., Yocum, D. R., York, D. G., Zehavi, I., and Zucker, D. B.: 2008, ApJS **175**, 297

Adelman-McCarthy, J. K. et al.: 2007, *VizieR Online Data Catalog* **2276**, 0

Alongi, M., Bertelli, G., Bressan, A., Chiosi, C., Fagotto, F., Greggio, L., and Nasi, E.: 1993, A&AS **97**, 851

Baldry, I. K., Glazebrook, K., and Driver, S. P.: 2008, MNRAS **388**, 945

Barber, T., Meiksin, A., and Murphy, T.: 2006, *preprint (astro-ph/0611053)*

Barrow, J. D., Juszkiewicz, R., and Sonoda, D. H.: 1985, MNRAS **213**, 917

Bell, E. F., McIntosh, D. H., Katz, N., and Weinberg, M. D.: 2003, ApJS **149**, 289

Bell, E. F., Papovich, C., Wolf, C., Le Floc'h, E., Caldwell, J. A. R., Barden, M., Egami, E., McIntosh, D. H., Meisenheimer, K., Pérez-González, P. G., Rieke, G. H., Rieke, M. J., Rigby, J. R., and Rix, H.-W.: 2005, ApJ **625**, 23

Bell, E. F., Zheng, X. Z., Papovich, C., Borch, A., Wolf, C., and Meisenheimer, K.: 2007, ApJ **663**, 834

Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Wollack, E., Wright, E. L., Barnes, C., Greason, M. R., Hill, R. S., Komatsu, E., Nolta, M. R., Odegard, N., Peiris, H. V., Verde, L., and Weiland, J. L.: 2003, ApJS **148**, 1

Benson, A. J., Bower, R. G., Frenk, C. S., Lacey, C. G., Baugh, C. M., and Cole, S.: 2003, ApJ **599**, 38

Binney, J. and Tremaine, S.: 2008, *Galactic Dynamics: Second Edition*, Galactic Dynamics: Second Edition, by James Binney and Scott Tremaine. ISBN 978-0-691-13026-2 (HB). Published by Princeton University Press, Princeton, NJ USA, 2008.

Blanton, M. R., Lupton, R. H., Schlegel, D. J., Strauss, M. A., Brinkmann, J., Fukugita, M., and Loveday, J.: 2005, ApJ **631**, 208

Borch, A., Meisenheimer, K., Bell, E. F., Rix, H.-W., Wolf, C., Dye, S., Kleinheinrich, M., Kovacs, Z., and Wisotzki, L.: 2006, A&A **453**, 869

Bower, R. G., Benson, A. J., Malbon, R., Helly, J. C., Frenk, C. S., Baugh, C. M., Cole, S., and Lacey, C. G.: 2006, MNRAS **370**, 645

Bressan, A., Fagotto, F., Bertelli, G., and Chiosi, C.: 1993, A&AS **100**, 647

Brinchmann, J., Charlot, S., White, S. D. M., Tremonti, C., Kauffmann, G., Heckman, T., and Brinkmann, J.: 2004, MNRAS **351**, 1151

Brinchmann, J. and Ellis, R. S.: 2000, ApJ **536**, L77

Bronder, T. J., Hook, I. M., Astier, P., Balam, D., Balland, C., Basa, S., Carlberg, R. G., Conley, A., Fouchez, D., Guy, J., Howell, D. A., Neill, J. D., Pain, R., Perrett, K., Pritchet, C. J., Regnault, N., Sullivan, M., Baumont, S., Fabbro, S., Filliol, M., Perlmutter, S., and Ripoche, P.: 2008, A&A **477**, 717

Bruzual, G. and Charlot, S.: 2003, MNRAS **344**, 1000

Bundy, K. et al.: 2006, ApJ **651**, 120

Carroll, S. M., Press, W. H., and Turner, E. L.: 1992, ARA&A **30**, 499

Castro, P. G.: 2003, Phys. Rev. D **67(12)**, 123001

Chabrier, G.: 2003, PASP **115**, 763

Charlot, S. and Fall, S. M.: 2000, ApJ **539**, 718

Cid Fernandes, R., Asari, N. V., Sodré, L., Stasińska, G., Mateus, A., Torres-Papaqui, J. P., and Schoenell, W.: 2007, MNRAS **375**, L16

Cid Fernandes, R., Gu, Q., Melnick, J., Terlevich, E., Terlevich, R., Kunth, D., Rodrigues Lacerda, R., and Joguet, B.: 2004, MNRAS **355**, 273

Clowe, D., Bradač, M., Gonzalez, A. H., Markevitch, M., Randall, S. W., Jones, C., and Zaritsky, D.: 2006, ApJ **648**, L109

Coelho, P., Bruzual, G., Charlot, S., Weiss, A., Barbuy, B., and Ferguson, J. W.: 2007, MNRAS **382**, 498

Cole, S., Norberg, P., Baugh, C. M., Frenk, C. S., Bland-Hawthorn, J., Bridges, T., Cannon, R., Colless, M., Collins, C., Couch, W., Cross, N., Dalton, G., De Propris, R., Driver, S. P., Efstathiou, G., Ellis, R. S., Glazebrook, K., Jackson, C., Lahav, O., Lewis, I., Lumsden, S., Maddox, S., Madgwick, D., Peacock, J. A., Peterson, B. A., Sutherland, W., and Taylor, K.: 2001, MNRAS **326**, 255

Cole, S., Percival, W. J., Peacock, J. A., Norberg, P., Baugh, C. M., Frenk, C. S., Baldry, I., Bland-Hawthorn, J., Bridges, T., Cannon, R., Colless, M., Collins, C., Couch, W., Cross, N. J. G., Dalton, G., Eke, V. R., De Propris, R., Driver, S. P., Efstathiou, G., Ellis, R. S., Glazebrook, K., Jackson, C., Jenkins, A., Lahav, O., Lewis, I., Lumsden, S., Maddox, S., Madgwick, D., Peterson, B. A., Sutherland, W., and Taylor, K.: 2005, MNRAS **362**, 505

Coles, P., Dineen, P., Earl, J., and Wright, D.: 2004, MNRAS **350**, 989

Coles, P. and Lucchin, F.: 1995, *Cosmology. The origin and evolution of cosmic structure*, Chichester: Wiley, —c1995

Colley, W. N. and Gott, J. R.: 2003, MNRAS **344**, 686

Copi, C. J., Huterer, D., and Starkman, G. D.: 2004, Phys. Rev. D **70(4)**, 043515

Cowie, L. L., Songaila, A., Hu, E. M., and Cohen, J. G.: 1996, AJ **112**, 839

Cruz, M., Cayón, L., Martínez-González, E., Vielva, P., and Jin, J.: 2007, ApJ **655**, 11

Cruz, M., Martínez-González, E., Vielva, P., and Cayón, L.: 2005, MNRAS **356**, 29

Dalal, N., Doré, O., Huterer, D., and Shirokov, A.: 2008, Phys. Rev. D **77(12)**, 123514

De Lucia, G., Springel, V., White, S. D. M., Croton, D., and Kauffmann, G.: 2006, MNRAS **366**, 499

Della Valle, M. and Panagia, N.: 2003, ApJ **587**, L71

Dennis, M. R. and Land, K.: 2008, MNRAS **383**, 424

Dilday, B., Kessler, R., Frieman, J. A., Holtzman, J., Marriner, J., Miknaitis, G., Nichol, R. C., Romani, R., Sako, M., Bassett, B., Becker, A., Cinabro, D., DeJongh, F., Depoy, D. L., Doi, M., Garnavich, P. M., Hogan, C. J., Jha, S., Konishi, K., Lampeitl, H., Marshall, J. L., McGinnis, D., Prieto, J. L., Riess, A. G., Richmond,

M. W., Schneider, D. P., Smith, M., Takanashi, N., Tokita, K., van der Heyden, K., Yasuda, N., Zheng, C., Barentine, J., Brewington, H., Choi, C., Crotts, A., Dembicky, J., Harvanek, M., Im, M., Ketzeback, W., Kleinman, S. J., Krzesiński, J., Long, D. C., Malanushenko, E., Malanushenko, V., McMillan, R. J., Nitta, A., Pan, K., Saurage, G., Snedden, S. A., Watters, S., Wheeler, J. C., and York, D.: 2008, ApJ **682**, 262

Dodelson, S.: 2003, *Modern cosmology*, Modern cosmology / Scott Dodelson. Amsterdam (Netherlands): Academic Press. ISBN 0-12-219141-2, 2003, XIII + 440 p.

Drory, N. and Alvarez, M.: 2008, ApJ **680**, 41

Efstathiou, G. and Bond, J. R.: 1986, MNRAS **218**, 103

Ellis, R. S., Sullivan, M., Nugent, P. E., Howell, D. A., Gal-Yam, A., Astier, P., Balam, D., Balland, C., Basa, S., Carlberg, R. G., Conley, A., Fouchez, D., Guy, J., Hardin, D., Hook, I., Pain, R., Perrett, K., Pritchet, C. J., and Regnault, N.: 2008, ApJ **674**, 51

Erb, D. K., Steidel, C. C., Shapley, A. E., Pettini, M., Reddy, N. A., and Adelberger, K. L.: 2006, ApJ **647**, 128

Eriksen, H. K., Banday, A. J., Górski, K. M., and Lilje, P. B.: 2004a, ApJ **612**, 633

Eriksen, H. K., Banday, A. J., Górski, K. M., and Lilje, P. B.: 2005, ApJ **622**, 58

Eriksen, H. K., Hansen, F. K., Banday, A. J., Górski, K. M., and Lilje, P. B.: 2004b, ApJ **605**, 14

Fagotto, F., Bressan, A., Bertelli, G., and Chiosi, C.: 1994a, A&AS **104**, 365

Fagotto, F., Bressan, A., Bertelli, G., and Chiosi, C.: 1994b, A&AS **105**, 29

Fioc, M. and Rocca-Volmerange, B.: 1999, *ArXiv Astrophysics e-prints*

Foley, R. J., Filippenko, A. V., Aguilera, C., Becker, A. C., Blondin, S., Challis, P., Clocchiatti, A., Covarrubias, R., Davis, T. M., Garnavich, P. M., Jha, S., Kirshner, R. P., Krisciunas, K., Leibundgut, B., Li, W., Matheson, T., Miceli, A., Miknaitis, G., Pignata, G., Rest, A., Riess, A. G., Schmidt, B. P., Smith, R. C., Sollerman, J., Spyromilio, J., Stubbs, C. W., Tonry, J. L., Suntzeff, N. B., Wood-Vasey, W. M., and Zenteno, A.: 2007, *ArXiv Astrophysics e-prints 0710.2338* 710

Fontana, A., Salimbeni, S., Grazian, A., Giallongo, E., Pentericci, L., Nonino, M., Fontanot, F., Menci, N., Monaco, P., Cristiani, S., Vanzella, E., de Santis, C., and Gallozzi, S.: 2006, A&A **459**, 745

Frieman, J. A., Bassett, B., Becker, A., Choi, C., Cinabro, D., DeJongh, F., Depoy, D. L., Dilday, B., Doi, M., Garnavich, P. M., Hogan, C. J., Holtzman, J., Im, M., Jha, S., Kessler, R., Konishi, K., Lampeitl, H., Marriner, J., Marshall, J. L., McGinnis, D., Miknaitis, G., Nichol, R. C., Prieto, J. L., Riess, A. G., Richmond, M. W., Romani, R., Sako, M., Schneider, D. P., Smith, M., Takanashi, N., Tokita, K., van der Heyden, K., Yasuda, N., Zheng, C., Adelman-McCarthy, J., Annis, J., Assef, R. J., Barentine, J., Bender, R., Blandford, R. D., Boroski, W. N., Bremer, M., Brewington, H., Collins, C. A., Crotts, A., Dembicky, J., Eastman, J., Edge, A., Edmondson, E., Elson, E., Eyler, M. E., Filippenko, A. V., Foley, R. J., Frank, S., Goobar, A., Gueth, T., Gunn, J. E., Harvanek, M., Hopp, U., Ihara, Y., Ivezić, Ž., Kahn, S., Kaplan, J., Kent, S., Ketzeback, W., Kleinman, S. J., Kollatschny, W., Kron, R. G., Krzesiński, J., Lamenti, D., Leloudas, G., Lin, H., Long, D. C., Lucey, J., Lupton, R. H., Malanushenko, E., Malanushenko, V., McMillan, R. J., Mendez, J., Morgan, C. W., Morokuma, T., Nitta, A., Ostman, L., Pan, K., Rockosi, C. M., Romer, A. K., Ruiz-Lapuente, P., Saurage, G., Schlesinger, K., Snedden, S. A., Sollerman, J., Stoughton, C., Stritzinger, M., Subba Rao, M., Tucker, D., Vaisanen, P., Watson, L. C., Watters, S., Wheeler, J. C., Yanny, B., and York, D.: 2008, AJ **135**, 338

Gallazzi, A., Charlot, S., Brinchmann, J., White, S. D. M., and Tremonti, C. A.: 2005, MNRAS **362**, 41

Girardi, L., Bressan, A., Chiosi, C., Bertelli, G., and Nasi, E.: 1996, A&AS **117**, 113

Glazebrook, K., Baldry, I. K., Blanton, M. R., Brinkmann, J., Connolly, A., Csabai, I., Fukugita, M., Ivezić, Ž., Loveday, J., Meiksin, A., Nichol, R., Peng, E., Schneider, D. P., SubbaRao, M., Tremonti, C., and York, D. G.: 2003a, ApJ **587**, 55

Glazebrook, K. et al.: 2003b, ApJ **587**, 55

Gold, B., Bennett, C. L., Hill, R. S., Hinshaw, G., Odegard, N., Page, L., Spergel, D. N., Weiland, J. L., Dunkley, J., Halpern, M., Jarosik, N., Kogut, A., Komatsu, E., Larson, D., Meyer, S. S., Nolta, M. R., Wollack, E., and Wright, E. L.: 2008, *ArXiv Astrophysics e-prints 0803.0715*

Gordon, K. D., Clayton, G. C., Misselt, K. A., Landolt, A. U., and Wolff, M. J.: 2003, ApJ **594**, 279

Gunn, J. E., Carr, M., Rockosi, C., Sekiguchi, M., Berry, K., Elms, B., de Haas, E., Ivezić, Ž., Knapp, G., Lupton, R., Pauls, G., Simcoe, R., Hirsch, R., Sanford, D., Wang, S., York, D., Harris, F., Annis, J., Bartozek, L., Boroski, W., Bakken, J., Haldeman, M., Kent, S., Holm, S., Holmgren, D., Petravick, D., Prosapio, A., Rechenmacher, R., Doi, M., Fukugita, M., Shimasaku, K., Okada, N., Hull, C., Siegmund, W., Mannery, E., Blouke, M., Heidtman, D., Schneider, D., Lucinio, R., and Brinkman, J.: 1998, AJ **116**, 3040

Gurzadyan, V. G., de Bernardis, P., de Troia, G., Bianco, C. L., Kashin, A. L., Kuloghlian, H., Masi, S., Piacentini, F., Polenta, G., and Yegorian, G.: 2005, *Modern Physics Letters A* **20**, 813

Guth, A. H.: 1981, Phys. Rev. D **23**, 347

Hamilton, A. J. S.: 1993, ApJ **417**, 19

Hamuy, M., Phillips, M. M., Suntzeff, N. B., Schommer, R. A., Maza, J., Smith, R. C., Lira, P., and Aviles, R.: 1996, AJ **112**, 2438

Heavens, A., Panter, B., Jimenez, R., and Dunlop, J.: 2004, Nature **428**, 625

Heavens, A. F. and Gupta, S.: 2001, *MNRAS* **324**, 960

Heavens, A. F., Jimenez, R., and Lahav, O.: 2000, MNRAS **317**, 965

Heavens, A. F. and Sheth, R. K.: 1999, *MNRAS* **310**, 1062

Hill, R. S., Weiland, J. L., Odegard, N., Wollack, E., Hinshaw, G., Larson, D., Bennett, C. L., Halpern, M., Page, L., Dunkley, J., Gold, B., Jarosik, N., Kogut, A., Limon, M., Nolta, M. R., Spergel, D. N., Tucker, G. S., and Wright, E. L.: 2008, *ArXiv Astrophysics e-prints 0803.0570*

Hinshaw, G., Barnes, C., Bennett, C. L., Greason, M. R., Halpern, M., Hill, R. S., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Odegard, N., Page, L., Spergel, D. N., Tucker, G. S., Weiland, J. L., Wollack, E., and Wright, E. L.: 2003a, ApJS **148**, 63

Hinshaw, G., Nolta, M. R., Bennett, C. L., Bean, R., Doré, O., Greason, M. R., Halpern, M., Hill, R. S., Jarosik, N., Kogut, A., Komatsu, E., Limon, M., Odegard, N., Meyer, S. S., Page, L., Peiris, H. V., Spergel, D. N., Tucker, G. S., Verde, L., Weiland, J. L., Wollack, E., and Wright, E. L.: 2007, ApJS **170**, 288

Hinshaw, G., Spergel, D. N., Verde, L., Hill, R. S., Meyer, S. S., Barnes, C., Bennett, C. L., Halpern, M., Jarosik, N., Kogut, A., Komatsu, E., Limon, M., Page, L., Tucker, G. S., Weiland, J. L., Wollack, E., and Wright, E. L.: 2003b, ApJS **148**, 135

Hook, I. M., Howell, D. A., Aldering, G., Amanullah, R., Burns, M. S., Conley, A., Deustua, S. E., Ellis, R., Fabbro, S., Fadeyev, V., Folatelli, G., Garavini, G., Gibbons, R., Goldhaber, G., Goobar, A., Groom, D. E., Kim, A. G., Knop, R. A., Kowalski, M., Lidman, C., Nobili, S., Nugent, P. E., Pain, R., Pennypacker, C. R., Perlmutter, S., Ruiz-Lapuente, P., Sainton, G., Schaefer, B. E., Smith, E., Spadafora, A. L., Stanishev, V., Thomas, R. C., Walton, N. A., Wang, L., and Wood-Vasey, W. M.: 2005, AJ **130**, 2788

Hopkins, A. M. and Beacom, J. F.: 2006, ApJ **651**, 142

Hopkins, A. M., Connolly, A. J., and Szalay, A. S.: 2000, AJ **120**, 2843

Howell, D. A.: 2001, ApJ **554**, L193

Howell, D. A., Sullivan, M., Conley, A., and Carlberg, R.: 2007, ApJ **667**, L37

Hu, W. and Dodelson, S.: 2002, ARA&A **40**, 171

Hu, W. T.: 1995, *Ph.D. Thesis*

Jaffe, T. R., Banday, A. J., Eriksen, H. K., Górski, K. M., and Hansen, F. K.: 2005, ApJ **629**, L1

Jarosik, N., Barnes, C., Greason, M. R., Hill, R. S., Nolta, M. R., Odegard, N., Weiland, J. L., Bean, R., Bennett, C. L., Doré, O., Halpern, M., Hinshaw, G., Kogut, A., Komatsu, E., Limon, M., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Wollack, E., and Wright, E. L.: 2007, ApJS **170**, 263

Jones, W. C., Ade, P. A. R., Bock, J. J., Bond, J. R., Borrill, J., Boscaleri, A., Cabella, P., Contaldi, C. R., Crill, B. P., de Bernardis, P., de Gasperis, G., de Oliveira-Costa, A., de Troia, G., di Stefano, G., Hivon, E., Jaffe, A. H., Kisner, T. S., Lange, A. E., MacTavish, C. J., Masi, S., Mauskopf, P. D., Melchiorri, A., Montroy, T. E., Natoli,

P., Netterfield, C. B., Pascale, E., Piacentini, F., Pogosyan, D., Polenta, G., Prunet, S., Ricciardi, S., Romeo, G., Ruhl, J. E., Santini, P., Tegmark, M., Veneziani, M., and Vittorio, N.: 2006, *New Astronomy Review* **50**, 945

Kauffmann, G., Colberg, J. M., Diaferio, A., and White, S. D. M.: 1999, MNRAS **303**, 188

Kauffmann, G. et al.: 2003, MNRAS **341**, 33

Kennicutt, Jr., R. C.: 1998, ApJ **498**, 541

Kereš, D., Katz, N., Weinberg, D. H., and Davé, R.: 2005, MNRAS **363**, 2

Komatsu, E., Dunkley, J., Nolta, M. R., Bennett, C. L., Gold, B., Hinshaw, G., Jarosik, N., Larson, D., Limon, M., Page, L., Spergel, D. N., Halpern, M., Hill, R. S., Kogut, A., Meyer, S. S., Tucker, G. S., Weiland, J. L., Wollack, E., and Wright, E. L.: 2008, *ArXiv Astrophysics e-prints 0803.0547*

Komatsu, E., Kogut, A., Nolta, M. R., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Limon, M., Meyer, S. S., Page, L., Spergel, D. N., Tucker, G. S., Verde, L., Wollack, E., and Wright, E. L.: 2003, ApJS **148**, 119

Komatsu, E. and Spergel, D. N.: 2001, Phys. Rev. D **63(6)**, 063002

Kroupa, P.: 2007, *ArXiv Astrophysics e-prints*

Lacey, C. and Cole, S.: 1993, MNRAS **262**, 627

Land, K. and Magueijo, J.: 2005a, MNRAS **357**, 994

Land, K. and Magueijo, J.: 2005b, *Physical Review Letters* **95(7)**, 071301

Land, K. and Magueijo, J.: 2005c, MNRAS **362**, 838

Larson, D. L. and Wandelt, B. D.: 2005, *ArXiv Astrophysics e-prints*

Lawson, C. and Hanson, R.: 1974, *Solving Least Squares Problems*, Prentice-Hall, Inc.

Le Borgne, J.-F., Bruzual, G., Pelló, R., Lançon, A., Rocca-Volmerange, B., Sanahuja, B., Schaerer, D., Soubiran, C., and Vílchez-Gómez, R.: 2003, A&A **402**, 433

Liguori, M., Matarrese, S., and Moscardini, L.: 2003, ApJ **597**, 57

Liu, X. and Zhang, S. N.: 2005, ApJ **633**, 542

Madau, P., Ferguson, H. C., Dickinson, M. E., Giavalisco, M., Steidel, C. C., and Fruchter, A.: 1996, MNRAS **283**, 1388

Mann, R. G., Saunders, W., and Taylor, A. N.: 1996, MNRAS **279**, 636

Mannucci, F., Della Valle, M., and Panagia, N.: 2006, MNRAS **370**, 773

Mannucci, F., Della Valle, M., Panagia, N., Cappellaro, E., Cresci, G., Maiolino, R., Petrosian, A., and Turatto, M.: 2005, A&A **433**, 807

Maoz, D.: 2008, MNRAS **384**, 267

Marigo, P., Girardi, L., Bressan, A., Groenewegen, M. A. T., Silva, L., and Granato, G. L.: 2008, A&A **482**, 883

Massey, R., Rhodes, J., Ellis, R., Scoville, N., Leauthaud, A., Finoguenov, A., Capak, P., Bacon, D., Aussel, H., Kneib, J.-P., Koekemoer, A., McCracken, H., Mobasher, B., Pires, S., Refregier, A., Sasaki, S., Starck, J.-L., Taniguchi, Y., Taylor, A., and Taylor, J.: 2007, Nature **445**, 286

Mathis, H., Charlot, S., and Brinchmann, J.: 2006, MNRAS **365**, 385

McEwen, J. D., Hobson, M. P., Lasenby, A. N., and Mortlock, D. J.: 2005, MNRAS **359**, 1583

McEwen, J. D., Hobson, M. P., Lasenby, A. N., and Mortlock, D. J.: 2006, MNRAS **369**, 1858

McEwen, J. D., Hobson, M. P., Lasenby, A. N., and Mortlock, D. J.: 2008, *ArXiv Astrophysics e-prints 0803.2157*

Milgrom, M.: 1983, ApJ **270**, 365

Mukherjee, P. and Wang, Y.: 2004, ApJ **613**, 51

Neill, J. D., Sullivan, M., Balam, D., Pritchet, C. J., Howell, D. A., Perrett, K., Astier, P., Aubourg, E., Basa, S., Carlberg, R. G., Conley, A., Fabbro, S., Fouchez, D., Guy, J., Hook, I., Pain, R., Palanque-Delabrouille, N., Regnault, N., Rich, J., Taillet, R., Aldering, G., Antilogus, P., Arsenijevic, V., Balland, C., Baumont, S., Bronder, J., Ellis, R. S., Filiol, M., Gonçalves, A. C., Hardin, D., Kowalski, M., Lidman, C., Lusset, V., Mouchet, M., Mourao, A., Perlmutter, S., Ripoche, P., Schlegel, D., and Tao, C.: 2006, AJ **132**, 1126

Noeske, K. G. et al.: 2007, *preprint (astro-ph/0703056)*

Nolta, M. R., Dunkley, J., Hill, R. S., Hinshaw, G., Komatsu, E., Larson, D., Page, L., Spergel, D. N., Bennett, C. L., Gold, B., Jarosik, N., Odegard, N., Weiland, J. L., Wollack, E., Halpern, M., Kogut, A., Limon, M., Meyer, S. S., Tucker, G. S., and Wright, E. L.: 2008, *ArXiv Astrophysics e-prints 0803.0593*

Ocvirk, P., Pichon, C., Lançon, A., and Thiébaut, E.: 2006, MNRAS **365**, 46

O'Donnell, J. E.: 1994, ApJ **422**, 158

Ostriker, J. P. and Vishniac, E. T.: 1986, Nature **322**, 804

Page, L., Barnes, C., Hinshaw, G., Spergel, D. N., Weiland, J. L., Wollack, E., Bennett, C. L., Halpern, M., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Tucker, G. S., and Wright, E. L.: 2003, ApJS **148**, 39

Paling, S. M.: 2005, *New Astronomy Review* **49**, 323

Panter, B., Heavens, A. F., and Jimenez, R.: 2003, MNRAS **343**, 1145

Panter, B., Jimenez, R., Heavens, A. F., and Charlot, S.: 2007, MNRAS **378**, 1550

Papovich, C., Moustakas, L. A., Dickinson, M., Le Floc'h, E., Rieke, G. H., Daddi, E., Alexander, D. M., Bauer, F., Brandt, W. N., Dahlen, T., Egami, E., Eisenhardt, P., Elbaz, D., Ferguson, H. C., Giavalisco, M., Lucas, R. A., Mobasher, B., Pérez-González, P. G., Stutz, A., Rieke, M. J., and Yan, H.: 2006, ApJ **640**, 92

Park, C.-G.: 2004, MNRAS **349**, 313

Peacock, J. A.: 1999, *Cosmological Physics*, Cosmological Physics, by John A. Peacock, pp. 704. ISBN 052141072X. Cambridge, UK: Cambridge University Press, January 1999.

Peacock, J. A. and Heavens, A. F.: 1990, MNRAS **243**, 133

Pearce, F. R., Jenkins, A., Frenk, C. S., White, S. D. M., Thomas, P. A., Couchman, H. M. P., Peacock, J. A., and Efstathiou, G.: 2001, MNRAS **326**, 649

Percival, W. J., Cole, S., Eisenstein, D. J., Nichol, R. C., Peacock, J. A., Pope, A. C., and Szalay, A. S.: 2007a, MNRAS **381**, 1053

Percival, W. J., Nichol, R. C., Eisenstein, D. J., Weinberg, D. H., Fukugita, M., Pope, A. C., Schneider, D. P., Szalay, A. S., Vogeley, M. S., Zehavi, I., Bahcall, N. A., Brinkmann, J., Connolly, A. J., Loveday, J., and Meiksin, A.: 2007b, ApJ **657**, 51

Perlmutter, S., Aldering, G., Goldhaber, G., Knop, R. A., Nugent, P., Castro, P. G., Deustua, S., Fabbro, S., Goobar, A., Groom, D. E., Hook, I. M., Kim, A. G., Kim, M. Y., Lee, J. C., Nunes, N. J., Pain, R., Pennypacker, C. R., Quimby, R., Lidman, C., Ellis, R. S., Irwin, M., McMahon, R. G., Ruiz-Lapuente, P., Walton, N., Schaefer, B., Boyle, B. J., Filippenko, A. V., Matheson, T., Fruchter, A. S., Panagia, N., Newberg, H. J. M., Couch, W. J., and The Supernova Cosmology Project: 1999, ApJ **517**, 565

Phillips, M. M.: 1993, ApJ **413**, L105

Pinsonneault, M. H. and Stanek, K. Z.: 2006, ApJ **639**, L67

Poelarends, A. J. T., Herwig, F., Langer, N., and Heger, A.: 2008, ApJ **675**, 614

Press, W. H. and Schechter, P.: 1974, ApJ **187**, 425

Pritchet, C. J., Howell, D. A., and Sullivan, M.: 2008, ApJ **683**, L25

Readhead, A. C. S., Mason, B. S., Contaldi, C. R., Pearson, T. J., Bond, J. R., Myers, S. T., Padin, S., Sievers, J. L., Cartwright, J. K., Shepherd, M. C., Pogosyan, D., Prunet, S., Altamirano, P., Bustos, R., Bronfman, L., Casassus, S., Holzapfel, W. L., May, J., Pen, U.-L., Torres, S., and Udomprasert, P. S.: 2004, ApJ **609**, 498

Rees, M. J., Sciama, D. W., and Setti, G.: 1968, Nature **217**, 326

Reichardt, C. L., Ade, P. A. R., Bock, J. J., Bond, J. R., Brevik, J. A., Contaldi, C. R., Daub, M. D., Dempsey, J. T., Goldstein, J. H., Holzapfel, W. L., Kuo, C. L., Lange, A. E., Lueker, M., Newcomb, M., Peterson, J. B., Ruhl, J., Runyan, M. C., and Staniszewski, Z.: 2008, *ArXiv Astrophysics e-prints 0801.1491*

Richmond, M. W., Filippenko, A. V., and Galisky, J.: 1998, PASP **110**, 553

Riess, A. G., Filippenko, A. V., Challis, P., Clocchiatti, A., Diercks, A., Garnavich, P. M., Gilliland, R. L., Hogan, C. J., Jha, S., Kirshner, R. P., Leibundgut, B., Phillips, M. M., Reiss, D., Schmidt, B. P., Schommer, R. A., Smith, R. C., Spyromilio, J., Stubbs, C., Suntzeff, N. B., and Tonry, J.: 1998, AJ **116**, 1009

Salpeter, E. E.: 1955, ApJ **121**, 161

Sarkar, D., Amblard, A., Cooray, A., and Holz, D. E.: 2008, *ArXiv Astrophysics e-prints 0806.3267*

Scannapieco, E. and Bildsten, L.: 2005, ApJ **629**, L85

Schlegel, D. J., Finkbeiner, D. P., and Davis, M.: 1998, ApJ **500**, 525

Sefusatti, E. and Komatsu, E.: 2007, Phys. Rev. D **76(8)**, 083004

Sheth, R. K., Jimenez, R., Panter, B., and Heavens, A. F.: 2006, ApJ **650**, L25

Sheth, R. K. and Tormen, G.: 2002, MNRAS **329**, 61

Slosar, A., Hirata, C., Seljak, U., Ho, S., and Padmanabhan, N.: 2008, *ArXiv Astrophysics e-prints 0805.3580*

Spergel, D. N., Bean, R., Doré, O., Nolta, M. R., Bennett, C. L., Dunkley, J., Hinshaw, G., Jarosik, N., Komatsu, E., Page, L., Peiris, H. V., Verde, L., Halpern, M., Hill, R. S., Kogut, A., Limon, M., Meyer, S. S., Odegard, N., Tucker, G. S., Weiland, J. L., Wollack, E., and Wright, E. L.: 2007, ApJS **170**, 377

Spergel, D. N. et al.: 2003, ApJS **148**, 175

Stark, P. B. and Parker, R. L.: 1995, *Computational Statistics* **10**, 143

Stoughton, C., Lupton, R. H., Bernardi, M., Blanton, M. R., Burles, S., Castander, F. J., Connolly, A. J., Eisenstein, D. J., Frieman, J. A., Hennessy, G. S., Hindsley, R. B., Ivezić, Ž., Kent, S., Kunszt, P. Z., Lee, B. C., Meiksin, A., Munn, J. A., Newberg, H. J., Nichol, R. C., Nicinski, T., Pier, J. R., Richards, G. T., Richmond, M. W., Schlegel, D. J., Smith, J. A., Strauss, M. A., SubbaRao, M., Szalay, A. S., Thakar, A. R., Tucker, D. L., Vanden Berk, D. E., Yanny, B., Adelman, J. K., Anderson, Jr., J. E., Anderson, S. F., Annis, J., Bahcall, N. A., Bakken, J. A., Bartelmann, M., Bastian, S., Bauer, A., Berman, E., Böhringer, H., Boroski, W. N., Bracker, S., Briegel, C., Briggs, J. W., Brinkmann, J., Brunner, R., Carey, L., Carr, M. A., Chen, B., Christian, D., Colestock, P. L., Crocker, J. H., Csabai, I., Czarapata, P. C., Dalcanton, J., Davidsen, A. F., Davis, J. E., Dehnen, W., Dodelson, S., Doi, M., Dombeck, T., Donahue, M., Ellman, N., Elms, B. R., Evans, M. L., Eyer, L., Fan, X., Federwitz, G. R., Friedman, S., Fukugita, M., Gal, R., Gillespie, B., Glazebrook, K., Gray, J., Grebel, E. K., Greenawalt, B., Greene, G., Gunn, J. E., de Haas, E.,

Haiman, Z., Haldeman, M., Hall, P. B., Hamabe, M., Hansen, B., Harris, F. H., Harris, H., Harvanek, M., Hawley, S. L., Hayes, J. J. E., Heckman, T. M., Helmi, A., Henden, A., Hogan, C. J., Hogg, D. W., Holmgren, D. J., Holtzman, J., Huang, C.-H., Hull, C., Ichikawa, S.-I., Ichikawa, T., Johnston, D. E., Kauffmann, G., Kim, R. S. J., Kimball, T., Kinney, E., Klaene, M., Kleinman, S. J., Klypin, A., Knapp, G. R., Korienek, J., Krolik, J., Kron, R. G., Krzesiński, J., Lamb, D. Q., Leger, R. F., Limmongkol, S., Lindenmeyer, C., Long, D. C., Loomis, C., Loveday, J., MacKinnon, B., Mannery, E. J., Mantsch, P. M., Margon, B., McGehee, P., McKay, T. A., McLean, B., Menou, K., Merelli, A., Mo, H. J., Monet, D. G., Nakamura, O., Narayanan, V. K., Nash, T., Neilsen, Jr., E. H., Newman, P. R., Nitta, A., Odenkirchen, M., Okada, N., Okamura, S., Ostriker, J. P., Owen, R., Pauls, A. G., Peoples, J., Peterson, R. S., Petravick, D., Pope, A., Pordes, R., Postman, M., Prosapio, A., Quinn, T. R., Rechenmacher, R., Rivetta, C. H., Rix, H.-W., Rockosi, C. M., Rosner, R., Ruthmansdorfer, K., Sandford, D., Schneider, D. P., Scranton, R., Sekiguchi, M., Sergey, G., Sheth, R., Shimasaku, K., Smee, S., Snedden, S. A., Stebbins, A., Stubbs, C., Szapudi, I., Szkody, P., Szokoly, G. P., Tabachnik, S., Tsvetanov, Z., Uomoto, A., Vogeley, M. S., Voges, W., Waddell, P., Walterbos, R., Wang, S.-i., Watanabe, M., Weinberg, D. H., White, R. L., White, S. D. M., Wilhite, B., Wolfe, D., Yasuda, N., York, D. G., Zehavi, I., and Zheng, W.: 2002, AJ **123**, 485

Strauss, M. A. et al.: 2002a, AJ **124**, 1810

Strauss, M. A., Weinberg, D. H., Lupton, R. H., Narayanan, V. K., Annis, J., Bernardi, M., Blanton, M., Burles, S., Connolly, A. J., Dalcanton, J., Doi, M., Eisenstein, D., Frieman, J. A., Fukugita, M., Gunn, J. E., Ivezić, Ž., Kent, S., Kim, R. S. J., Knapp, G. R., Kron, R. G., Munn, J. A., Newberg, H. J., Nichol, R. C., Okamura, S., Quinn, T. R., Richmond, M. W., Schlegel, D. J., Shimasaku, K., SubbaRao, M., Szalay, A. S., Vanden Berk, D., Vogeley, M. S., Yanny, B., Yasuda, N., York, D. G., and Zehavi, I.: 2002b, AJ **124**, 1810

Sullivan, M., Le Borgne, D., Pritchet, C. J., Hodsman, A., Neill, J. D., Howell, D. A., Carlberg, R. G., Astier, P., Aubourg, E., Balam, D., Basa, S., Conley, A., Fabbro, S., Fouchez, D., Guy, J., Hook, I., Pain, R., Palanque-Delabrouille, N., Perrett, K., Regnault, N., Rich, J., Taillet, R., Baumont, S., Bronder, J., Ellis, R. S., Filiol, M., Lusset, V., Perlmutter, S., Ripoche, P., and Tao, C.: 2006, ApJ **648**, 868

Sunyaev, R. A. and Zeldovich, Y. B.: 1972, A&A **20**, 189

Tegmark, M., Blanton, M. R., Strauss, M. A., Hoyle, F., Schlegel, D., Scoccimarro, R., Vogeley, M. S., Weinberg, D. H., Zehavi, I., Berlind, A., Budavari, T., Connolly, A., Eisenstein, D. J., Finkbeiner, D., Frieman, J. A., Gunn, J. E., Hamilton, A. J. S., Hui, L., Jain, B., Johnston, D., Kent, S., Lin, H., Nakajima, R., Nichol, R. C., Ostriker, J. P., Pope, A., Scranton, R., Seljak, U., Sheth, R. K., Stebbins, A., Szalay, A. S., Szapudi, I., Verde, L., Xu, Y., Annis, J., Bahcall, N. A., Brinkmann, J., Burles, S., Castander, F. J., Csabai, I., Loveday, J., Doi, M., Fukugita, M., Gott, J. R. I., Hennessy, G., Hogg, D. W., Ivezić, Ž., Knapp, G. R., Lamb, D. Q., Lee, B. C., Lupton, R. H., McKay, T. A., Kunszt, P., Munn, J. A., O'Connell, L., Peoples, J., Pier, J. R., Richmond, M., Rockosi, C., Schneider, D. P., Stoughton, C., Tucker, D. L., Vanden Berk, D. E., Yanny, B., and York, D. G.: 2004, ApJ **606**, 702

Tegmark, M., de Oliveira-Costa, A., and Hamilton, A. J.: 2003, Phys. Rev. D **68(12)**, 123523

Thomas, D., Maraston, C., and Bender, R.: 2003, MNRAS **339**, 897

Thomas, D., Maraston, C., Bender, R., and Mendes de Oliveira, C.: 2005, ApJ **621**, 673

Tojeiro, R., Castro, P. G., Heavens, A. F., and Gupta, S.: 2006, MNRAS **365**, 265

Totani, T., Morokuma, T., Oda, T., Doi, M., and Yasuda, N.: 2008, *ArXiv Astrophysics e-prints 0804.0909*

Tremonti, C. A. et al.: 2004, ApJ **613**, 898

van den Bergh, S., Li, W., and Filippenko, A. V.: 2005, PASP **117**, 773

Verma, A., Lehnert, M. D., Foerster Schreiber, N. M., Bremer, M. N., and Douglas, L.: 2007, *preprint (astro-ph/0701725)*

Vielva, P., Martínez-González, E., Barreiro, R. B., Sanz, J. L., and Cayón, L.: 2004, ApJ **609**, 22

Wall, J. V. and Jenkins, C.: 2003, *Practical statistics fpr Astronomers*, Cambridge University Press

Weinberg, D. H., Davé, R., Katz, N., and Hernquist, L.: 2004, ApJ **601**, 1

White, S. D. M. and Rees, M. J.: 1978, MNRAS **183**, 341

Wiaux, Y., Vielva, P., Barreiro, R. B., Martínez-González, E., and Vandergheynst, P.: 2008, MNRAS **385**, 939

Wilkins, S. M., Trentham, N., and Hopkins, A. M.: 2008, MNRAS **385**, 687

Wood-Vasey, W. M., Miknaitis, G., Stubbs, C. W., Jha, S., Riess, A. G., Garnavich, P. M., Kirshner, R. P., Aguilera, C., Becker, A. C., Blackman, J. W., Blondin, S., Challis, P., Clocchiatti, A., Conley, A., Covarrubias, R., Davis, T. M., Filippenko, A. V., Foley, R. J., Garg, A., Hicken, M., Krisciunas, K., Leibundgut, B., Li, W., Matheson, T., Miceli, A., Narayan, G., Pignata, G., Prieto, J. L., Rest, A., Salvo, M. E., Schmidt, B. P., Smith, R. C., Sollerman, J., Spyromilio, J., Tonry, J. L., Suntzeff, N. B., and Zenteno, A.: 2007, ApJ **666**, 694

Worthey, G.: 1994, ApJS **95**, 107

Wright, E. L., Chen, X., Odegard, N., Bennett, C. L., Hill, R. S., Hinshaw, G., Jarosik, N., Komatsu, E., Nolta, M. R., Page, L., Spergel, D. N., Weiland, J. L., Wollack, E., Dunkley, J., Gold, B., Halpern, M., Kogut, A., Larson, D., Limon, M., Meyer, S. S., and Tucker, G. S.: 2008, *ArXiv Astrophysics e-prints 0803.0577*

Yadav, A. P. S. and Wandelt, B. D.: 2008, *Physical Review Letters* **100(18)**, 181301

York, D. G., Adelman, J., Anderson, Jr., J. E., Anderson, S. F., Annis, J., Bahcall, N. A., Bakken, J. A., Barkhouser, R., Bastian, S., Berman, E., Boroski, W. N., Bracker, S., Briegel, C., Briggs, J. W., Brinkmann, J., Brunner, R., Burles, S., Carey, L., Carr, M. A., Castander, F. J., Chen, B., Colestock, P. L., Connolly, A. J., Crocker, J. H., Csabai, I., Czarapata, P. C., Davis, J. E., Doi, M., Dombeck, T., Eisenstein, D., Ellman, N., Elms, B. R., Evans, M. L., Fan, X., Federwitz, G. R., Fiscelli, L., Friedman, S., Frieman, J. A., Fukugita, M., Gillespie, B., Gunn, J. E., Gurbani, V. K., de Haas, E., Haldeman, M., Harris, F. H., Hayes, J., Heckman, T. M., Hennessy, G. S., Hindsley, R. B., Holm, S., Holmgren, D. J., Huang, C.-h., Hull, C., Husby, D., Ichikawa, S.-I., Ichikawa, T., Ivezić, Ž., Kent, S., Kim, R. S. J., Kinney, E., Klaene, M., Kleinman, A. N., Kleinman, S., Knapp, G. R., Korienek, J., Kron, R. G., Kunszt, P. Z., Lamb, D. Q., Lee, B., Leger, R. F., Limmongkol, S., Lindenmeyer, C., Long, D. C., Loomis, C., Loveday, J., Lucinio, R., Lupton, R. H., MacKinnon, B., Mannery, E. J., Mantsch, P. M., Margon, B., McGehee, P., McKay, T. A., Meiksin, A., Merelli, A., Monet, D. G., Munn, J. A., Narayanan, V. K., Nash,

T., Neilsen, E., Neswold, R., Newberg, H. J., Nichol, R. C., Nicinski, T., Nonino, M., Okada, N., Okamura, S., Ostriker, J. P., Owen, R., Pauls, A. G., Peoples, J., Peterson, R. L., Petravick, D., Pier, J. R., Pope, A., Pordes, R., Prosapio, A., Rechenmacher, R., Quinn, T. R., Richards, G. T., Richmond, M. W., Rivetta, C. H., Rockosi, C. M., Ruthmansdorfer, K., Sandford, D., Schlegel, D. J., Schneider, D. P., Sekiguchi, M., Sergey, G., Shimasaku, K., Siegmund, W. A., Smee, S., Smith, J. A., Snedden, S., Stone, R., Stoughton, C., Strauss, M. A., Stubbs, C., SubbaRao, M., Szalay, A. S., Szapudi, I., Szokoly, G. P., Thakar, A. R., Tremonti, C., Tucker, D. L., Uomoto, A., Vanden Berk, D., Vogeley, M. S., Waddell, P., Wang, S.-i., Watanabe, M., Weinberg, D. H., Yanny, B., and Yasuda, N.: 2000a, AJ **120**, 1579

York, D. G. et al.: 2000b, AJ **120**, 1579

Zheng, X. Z., Bell, E. F., Papovich, C., Wolf, C., Meisenheimer, K., Rix, H.-W., Rieke, G. H., and Somerville, R.: 2007, ApJ **661**, L41

Zwicky, F.: 1933, *Helvetica Physica Acta* **6**, 110